ONTAP 9.16.1 commands

ONTAP 9.16.1 commands

NetApp
December 18, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-cli/index.html on December 18, 2024.
Always check docs.netapp.com for the latest.

Table of Contents

ONTAP 9.16.1 commands . . . 1
exit . . . 2
  Description . . . 2
  Examples . . . 2
history . . . 3
  Description . . . 3
  Examples . . . 3
  Related Links . . . 3
man . . . 4
  Description . . . 4
  Parameters . . . 4
  Examples . . . 4
redo . . . 5
  Description . . . 5
  Parameters . . . 5
  Examples . . . 5
  Related Links . . . 5
rows . . . 6
  Description . . . 6
  Parameters . . . 6
  Examples . . . 6
  Related Links . . . 6
set . . . 7
  Description . . . 7
  Parameters . . . 7
  Examples . . . 9
  Related Links . . . 10
top . . . 11
  Description . . . 11
  Examples . . . 11
up . . . 12
  Description . . . 12
  Examples . . . 12
application commands . . . 13
  application provisioning commands . . . 13
autobalance commands . . . 14
  autobalance aggregate commands . . . 14
cluster commands . . . 22
  cluster add-node-status-clear-failed . . . 22
  cluster add-node-status . . . 23
  cluster add-node . . . 24
  cluster create . . . 25
  cluster join . . . 27

cluster modify . . . 28
cluster ping-cluster . . . 29
cluster remove-node . . . 30
cluster setup . . . 31
cluster show . . . 41
cluster agent commands . . . 42
cluster contact-info commands . . . 49
cluster controller-replacement commands . . . 51
cluster date commands . . . 55
cluster ha commands . . . 60
cluster identity commands . . . 61
cluster image commands . . . 62
cluster kernel-service commands . . . 91
cluster log-forwarding commands . . . 95
cluster peer commands . . . 101
cluster quorum-service commands . . . 136
cluster ring commands . . . 138
cluster statistics commands . . . 139
cluster time-service commands . . . 141
event commands . . . 154
event catalog commands . . . 154
event config commands . . . 157
event filter commands . . . 162
event log commands . . . 187
event notification commands . . . 191
event role-config commands . . . 209
event status commands . . . 214
job commands . . . 218
job delete . . . 218
job pause . . . 218
job resume . . . 219
job show-bynode . . . 220
job show-cluster . . . 221
job show-completed . . . 224
job show . . . 226
job stop . . . 229
job unclaim . . . 230
job watch-progress . . . 230
job history commands . . . 231
job initstate commands . . . 234
job private commands . . . 236
job schedule commands . . . 245
lun commands . . . 258
lun convert-from-namespace . . . 258
lun create . . . 258

lun delete . . . 262
lun maxsize . . . 263
lun modify . . . 266
lun move-in-volume . . . 269
lun resize . . . 271
lun show . . . 273
lun bind commands . . . 280
lun copy commands . . . 283
lun igroup commands . . . 290
lun import commands . . . 304
lun mapping commands . . . 315
lun move commands . . . 323
lun persistent-reservation commands . . . 329
lun portset commands . . . 332
lun transition commands . . . 337
metrocluster commands . . . 343
metrocluster configure . . . 343
metrocluster heal . . . 347
metrocluster modify . . . 347
metrocluster remove-dr-group . . . 349
metrocluster show . . . 350
metrocluster switchback . . . 352
metrocluster switchover . . . 353
metrocluster check commands . . . 355
metrocluster config-replication commands . . . 387
metrocluster configuration-settings commands . . . 390
metrocluster interconnect commands . . . 425
metrocluster node commands . . . 435
metrocluster operation commands . . . 440
metrocluster transition commands . . . 443
metrocluster vserver commands . . . 445
network commands . . . 449
network ping . . . 449
network ping6 . . . 450
network test-path . . . 452
network traceroute . . . 453
network traceroute6 . . . 455
network arp commands . . . 457
network bgp commands . . . 461
network cloud commands . . . 471
network connections commands . . . 473
network device-discovery commands . . . 484
network fcp commands . . . 486
network interface commands . . . 497
network ipspace commands . . . 546

network ndp commands . . . 549
network options commands . . . 557
network port commands . . . 568
network qos-marking commands . . . 596
network route commands . . . 598
network subnet commands . . . 604
network test-link commands . . . 610
network tuning commands . . . 614
protection-type commands . . . 620
protection-type show . . . 620
qos commands . . . 622
qos adaptive-policy-group commands . . . 622
qos policy-group commands . . . 628
qos settings commands . . . 634
qos statistics commands . . . 637
qos workload commands . . . 678
san commands . . . 681
san config commands . . . 681
security commands . . . 682
security snmpusers . . . 682
security anti-ransomware commands . . . 683
security audit commands . . . 706
security certificate commands . . . 710
security config commands . . . 752
security cryptomod-fips commands . . . 767
security dynamic-authorization commands . . . 768
security ipsec commands . . . 787
security key-manager commands . . . 807
security login commands . . . 876
security multi-admin-verify commands . . . 956
security oauth2 commands . . . 974
security protocol commands . . . 982
security saml-sp commands . . . 985
security session commands . . . 990
security ssh commands . . . 1030
security ssl commands . . . 1037
security tpm commands . . . 1042
security webauthn commands . . . 1043
snaplock commands . . . 1049
snaplock compliance-clock commands . . . 1049
snaplock event-retention commands . . . 1051
snaplock legal-hold commands . . . 1058
snaplock log commands . . . 1065
snapmirror commands . . . 1071
snapmirror abort . . . 1071

snapmirror break . . . 1074
snapmirror create . . . 1077
snapmirror delete . . . 1085
snapmirror flexgroup-epuuid-prefix-prepare-to-downgrade . . . 1088
snapmirror initialize-ls-set . . . 1089
snapmirror initialize . . . 1091
snapmirror list-destinations . . . 1097
snapmirror modify . . . 1103
snapmirror promote . . . 1107
snapmirror protect . . . 1109
snapmirror quiesce . . . 1113
snapmirror release . . . 1116
snapmirror restore . . . 1118

ONTAP 9.16.1 commands

exit

Quit the CLI session

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

Description

The `exit` command ends the current CLI session.

Examples

The following example ends the current CLI session:

```
cluster1::> exit
Goodbye
```

history

Show the history of commands for this CLI session

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

Description

The `history` command displays the command history of the current CLI session. A numeric ID precedes each command. Use this number with the redo command to re-execute that history item.

Examples

The following example displays the command history of the current CLI session:

```
cluster1::> history
   1  vserver show
   2  man volume show
   3  volume delete -vserver vs0 -volume temporary2
   4  volume modify { -volume temp* } -state offline
cluster1::> redo 3
```

Related Links

- redo

man

Display the online manual pages

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

Description

The `man` command displays the manual page of the command you specify. If you do not specify a command, command displays the man page index.

Parameters

`[<text>]` - Valid CLI command

The command for which you'd like to see the manual page. The syntax of the command is the same as the command itself. The `man` command supports abbreviations and tab completion of the command name.

Examples

The following example displays the manual page for the *`storage aggregate create`* command.

```
cluster1::> man sto aggr cre
```

That example could also have been fully specified as:

```
cluster1::> man storage aggregate create
```

redo

Execute a previous command

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

Description

The `redo` command re-executes a command that has been executed previously in the current CLI session. Specify a previously run command using:

- A string that matches part of a previous command. For example, if the only *`volume`* command you have run is *`volume show`* , enter `redo``vol` to re-execute the command.
- The numeric ID of a previous command, as listed by the history command. For example, enter `redo``4` to re-execute the fourth command in the history list.
- A negative offset from the end of the history list. For example, enter `redo`-2` to re-execute the command that you ran two commands ago.

Parameters

`[<text>]` - String, Event Number, or Negative Offset

Use this parameter to specify a string, a numeric ID from the command history, or a negative number that identifies the command to be re-executed.

Examples

The following example re-executes command number 10 in the command history:

```
cluster1::> redo 10
```

Related Links

- history

rows

Show/Set the rows for the CLI session

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

Description

The `rows` command displays the number of rows that can be displayed in the current CLI session before the interface pauses output. If you do not set this value, it adjusts automatically based on the actual height of your terminal. If the actual height is undefined, the default number of rows is 24.

Specify a number to set the number of rows that can be displayed. Setting this value manually disables auto-adjustment. Specify zero (`0`) to disable pausing.

You can also set this value using the set *`-rows`* command.

Parameters

`[<integer>]` - Number of Rows the Screen Can Display

Use this parameter to specify the number of rows your terminal can display.

Examples

The following example displays the current number of rows, then resets the number of rows to 48:

```
cluster1::> rows
  36

cluster1::> rows 48
```

Related Links

- set

set

Display/Set CLI session settings

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

Description

The `set` command changes attributes of the user interface.

Parameters

`[-privilege <PrivilegeLevel>]` - Privilege Level

Use this parameter to specify the privilege level of the command session. Possible values are

- admin - Used for routine system management commands
- advanced - Used for infrequent, dangerous, or complicated commands
- diagnostic - Used for detailed diagnostic commands that are used only by support personnel

`[-confirmations {on|off}]` - Confirmation Messages

Use this parameter with the value `on` to specify that the interface prompt for confirmation before executing potentially dangerous commands. Use this parameter with the value `off` to specify that the interface not prompt for confirmation, even before potentially dangerous commands execute. The default setting is `on` .

`[-showallfields {true|false}]` - Show All Fields

Use this parameter with the value `true` to specify that the interface display all field columns when displaying tabular output. Use this parameter with the value `false` to specify that the interface display only selected columns. The default setting is `false` .

`[-showseparator <text>]` - Show Separator

Use this parameter to specify the characters to use as the field separator. The field separator is used between field columns when `-showallfields` is set to "true". The separator can be from one to three characters in length. When specifying the separator, enclose it in quotation marks (`"`). Set the separator to one or more spaces to disable this feature.

`[-active-help {true|false}]` - Active Help

Use this parameter with the value `true` to specify that pressing the question mark (`?`) key is sufficient to execute a help request. Use this parameter with the value `false` to specify that you must press the Return key after the question mark key to execute a help request. The default setting is `true` .

`[-units {auto|raw|B|KB|MB|GB|TB|PB}]` - Data Units

Use this parameter to specify the default units used when reporting data sizes. Possible values are:

- `auto` - Auto-scale data size for human-readable output
- `raw` - Bytes without unit designation
- `B` - Bytes

- `KB` - Kilobytes (1024 bytes, aka kibibytes)
- `MB` - Megabytes (KB x 1024, aka mebibytes)
- `GB` - Gigabytes (MB x 1024, aka gibibytes)
- `TB` - Terabytes (GB x 1024, aka tebibytes)
- `PB` - Petabytes (TB x 1024, aka pebibytes)

The default setting is `auto` .

`[-rows <integer>]` - Pagination Rows ('0' disables)

Use this parameter to specify the number of rows that can be displayed in the current CLI session before the interface pauses output. If you do not set this value, it adjusts automatically based on the actual height of your terminal. If the actual height is undefined, the default number of rows is 24.

Setting this value manually disables auto-adjustment. Specify zero (`0`) to disable pausing.

You can also set this value using the rows command.

`[-vserver <text>]` - Default Vserver

Use this parameter to specify the name of the Vserver to use as the default value for the `-vserver` parameter of commands.

> Vserverized commands that only have a single required parameter, which is the <userinput>-vserver<userinput>, allow the Vserver to be specified positionally, without <userinput>-vserver<userinput> preceding it. Due to this, care must be taken when using CLI commands that do not require the <userinput>-vserver<userinput> parameter. For example, using the "vserver nfs delete **" command will ignore the "set -vserver" value as the parser considers the ""** to be the Vserver.

`[-node <text>]` - Default Node

Use this parameter to specify the name of the node to use as the default value for the `-node` parameter of commands.

`[-stop-on-error {true|false}]` - Stop On Error

Use this parameter with the value `true` to specify that continuing commands should stop if they encounter an error. Use this parameter with the value `false` to specify that continuing commands should continue if they encounter an error.

`[-prompt-timestamp {above|inline|none}]` - Display Prompt Timestamp

Print the current date and time as a part of the prompt. The possible values are

- *`above`* - print the timestamp using the system timestamp format on the line above the remainder of the prompt.
- *`inline`* - print the timestamp using the system timestamp format at the beginning of the line with the remainder of the prompt.
- *`none`* - do not print the timestamp.

The default value is *`none`* .

Examples

The following example sets the privilege level to advanced.

```
cluster1::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only
when
         directed to do so by NetApp personnel.
Do you wish to continue? (y or n): y

cluster1::*>
```

The following examples cause all columns to be shown in output rows, with a comma used as the field separator.

```
cluster1::> set -showallfields true

cluster1::> set -showseparator ","

cluster1::> network port show
node,port,role,link,mtu,autonegotiate-admin,autonegotiate-oper,duplex-
admin,duplex-oper,speed-admin,speed-oper,flowcontrol-admin,flowcontrol-
oper,mac,up-admin,type,ifgrp-node,ifgrp-port,ifgrp-distr-func,ifgrp-
mode,vlan-node,vlan-port,vlan-tag,
Node,Port,Role,Link,MTU,Auto-Negotiation Administrative,Auto-Negotiation
Operational,Duplex Mode Administrative,Duplex Mode Operational,Speed
Administrative,Speed Operational,Flow Control Administrative,Flow Control
Operational,MAC Address,Up Administrative,Port Type,Interface Group Parent
Node,Interface Group Parent Port,Distribution,Create Policy,Parent VLAN
Node,Parent VLAN Port,VLAN Tag,
node1,e0a,cluster,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29
:90:20:e9,true,physical,-,-,-,-,-,-,-,
node1,e0b,cluster,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29
:90:20:f3,true,physical,-,-,-,-,-,-,-,
node1,e0c,data,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29:90
:20:fd,true,physical,-,-,-,-,-,-,-,
node1,e0d,data,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29:90
:20:07,true,physical,-,-,-,-,-,-,-,
node2,e0a,cluster,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29
:2e:b6:62,true,physical,-,-,-,-,-,-,-,
node2,e0b,cluster,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29
:2e:b6:6c,true,physical,-,-,-,-,-,-,-,
node2,e0c,data,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29:2e
:b6:76,true,physical,-,-,-,-,-,-,-,
node2,e0d,data,up,1500,true,true,full,full,auto,1000,full,none,00:0c:29:2e
:b6:80,true,physical,-,-,-,-,-,-,-,
```

The following example shows how to create a prompt with a timestamp.

```
cluster1::> set -prompt-timestamp above
[2/25/2016 16:38:38]
 cluster1::>
```

Related Links

- rows

top

Go to the top-level directory

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

Description

The `top` command changes the current working directory of the command prompt to the top-level command directory.

Examples

The following example returns the command prompt from the `storage aggregate` directory to the top-level directory:

```
cluster1::storage aggregate> top

cluster1::>
```

up

Go up one directory

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

Description

The `up` command, which can also be specified as two dots (`..`), changes the current working directory of the command prompt to the directory that is up one level in the command hierarchy.

Examples

The following example takes the command prompt up one level from the `storage aggregate` directory:

```
cluster1::storage aggregate> up

cluster1::storage>
```

application commands

application provisioning commands

application provisioning config modify

Modify options for application provisioning

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

Description

This command modifies the options for application provisioning operations.

Parameters

`[-is-mixed-storage-services-allowed {true|false}]` - Is Mixed Storage Services Allowed (privilege: advanced)

Specifies whether mixed cost storage services are allowed for provisioning placement. If the value of this parameter is *`false`* , only the aggregates closest to the performance requirements of the storage service are used. If the value of this parameter is *`true`* , all aggregates with sufficient performance are considered. The initial value for option is *`false`* .

Examples

```
cluster1::*> application provisioning config modify -is-mixed-storage
-services-allowed true
```

Enables the use of mixed storage services for provisioning placement.

application provisioning config show

Display options for application provisioning

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

Description

This command displays options for application provisioning.

Examples

```
cluster1::*> application provisioning config show
Is Mixed Storage Services Allowed: false
```

autobalance commands

autobalance aggregate commands

autobalance aggregate show-aggregate-state

Display the Auto Balance Aggregate state for an aggregate

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

Description

The `autobalance aggregate show-aggregate-state` command displays information about an aggregate state that is considered by the Auto Balance Aggregate feature.

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>|local}]` - Node Name (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with a node that matches the specified value.

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

If this parameter is specified, the display will be limited to only that aggregate with a name that matches the specified value.

`[-total-size {<integer>[KB|MB|GB|TB|PB]}]` - Total Size of the Aggregate (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with a total-size that matches the specified value.

`[-used-size {<integer>[KB|MB|GB|TB|PB]}]` - Used Size of the Aggregate (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with a used-size that matches the specified value.

`[-aggregate-unbalanced-threshold {<integer>[KB|MB|GB|TB|PB]}]` - Threshold When Aggregate Is Considered Unbalanced (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with a threshold that matches the specified value.

`[-outgoing-size {<integer>[KB|MB|GB|TB|PB]}]` - Size of Outgoing Volumes in the Aggregate (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with an outgoing-size that matches the specified value. Outgoing size will be equal to the total size of the volumes that move away from each one of those aggregate.

[-incoming-size {<integer>[KB|MB|GB|TB|PB]}] - Size of Incoming Volumes in the Aggregate (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with an incoming-size that matches the specified value. Incoming size will be equal to the total size of the volumes that move towards to each one of those aggregates.

[-raidtype {raid_tec|raid_dp|raid4|raid_ep}] - RAID Type (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with a raidtype that matches the specified value.

[-home-cluster <UUID>] - Home Cluster ID (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with a home-cluster ID that matches the specified value.

[-is-hybrid {true|false}] - Aggregate Is a Hybrid (privilege: advanced)

If this parameter is specified as true, the display will be limited to only hybrid aggregates. If the parameter is specified as false, the display will be limited to only non-hybrid aggregates.

[-is-incoming-volume-thin {true|false}] - An Incoming Volume Is Thin (privilege: advanced)

When you use thin provisioning for a volume, it can run out of space even if it has not yet consumed its nominal size and you should carefully monitor space utilization to avoid unexpected errors due to the volume running out of space. If this parameter is specified as true, the display will be limited to only those aggregates which are the target of a move of thin volume. If the parameter is specified as false, the display will be limited to only those aggregates which are not the target of a move of thin volume.

[-is-balanceable {true|false}] - Is Balanceable (privilege: advanced)

If this parameter is specified as true, the display will be limited to only balanceable aggregates. If the parameter is specified as false, the display will be limited to only non-balanceable aggregates.

[-is-move-target {true|false}] - Aggregate Is a Volume Move Target (privilege: advanced)

If this parameter is specified as true, the display will be limited to only those aggregates which are target of a volume move. If the parameter is specified as false, the display will be limited to only those aggregates which are not the target of a volume move.

[-attributes <text>,…] - Aggregate Attributes (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates with attributes that matches the specified values.

[-aggregate-available-threshold {<integer>[KB|MB|GB|TB|PB]}] - Threshold When Aggregate Is Considered Balanced (privilege: advanced)

If this parameter is specified, the display will be limited to only those aggregates which meet the specified threshold to be considered as balanced.

Examples

The following example displays information about the state for all aggregates in the cluster.

```
cluster1::*> autobalance aggregate show-aggregate-state
                          Aggregate: aggr0
                         Total Size: 4.78GB
                          Used Size: 4.56GB
                       Outgoing Size: 0B
                       Incoming Size: 0B
     Aggregate Used Space Threshold: 3.34GB
Aggregate Available Space Threshold: 1.91GB
                          RAID Type: raid_dp
                     Home Cluster ID: edf0379b-16da-11e6-aa3c-0050568558c2
                         Attributes: CFO
                                     Excluded
                                     Mroot
Aggregate: aggr_1
                         Total Size: 12.61GB
                          Used Size: 111.6MB
                       Outgoing Size: 0B
                       Incoming Size: 0B
     Aggregate Used Space Threshold: 8.83GB
Aggregate Available Space Threshold: 5.04GB
                          RAID Type: raid4
                     Home Cluster ID: edf0379b-16da-11e6-aa3c-0050568558c2
                         Attributes: Excluded
```

The following example displays information about all entries of the aggregate state, for all aggregates in the cluster.

```
cluster1::*> autobalance aggregate show-aggregate-state -instance
                                          Node Name: cluster-1-01
                             Name of the Aggregate: aggr0
                       Total Size of the Aggregate: 4.78GB
                        Used Size of the Aggregate: 4.56GB
Threshold When Aggregate Is Considered Unbalanced: 3.34GB
        Size of Outgoing Volumes in the Aggregate: 0B
        Size of Incoming Volumes in the Aggregate: 0B
                                          RAID Type: raid_dp
                                     Home Cluster ID: edf0379b-16da-11e6-
aa3c-0050568558c2
                             Aggregate Is a Hybrid: false
                        An Incoming Volume Is Thin: false
                                    Is Balanceable: false
                 Aggregate Is a Volume Move Target: false
                               Aggregate Attributes: CFO
                                                     Excluded
                                                     Mroot
  Threshold When Aggregate Is Considered Balanced: 1.91GB
Node Name: cluster-1-01
                             Name of the Aggregate: aggr_1
                       Total Size of the Aggregate: 12.61GB
                        Used Size of the Aggregate: 111.6MB
Threshold When Aggregate Is Considered Unbalanced: 8.83GB
        Size of Outgoing Volumes in the Aggregate: 0B
        Size of Incoming Volumes in the Aggregate: 0B
                                          RAID Type: raid4
                                     Home Cluster ID: edf0379b-16da-11e6-
aa3c-0050568558c2
                             Aggregate Is a Hybrid: false
                        An Incoming Volume Is Thin: false
                                    Is Balanceable: false
                 Aggregate Is a Volume Move Target: false
                               Aggregate Attributes: Excluded
  Threshold When Aggregate Is Considered Balanced: 5.04GB
```

autobalance aggregate show-unbalanced-volume-state

Display the Auto Balance Aggregate state for a volume

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

Description

The `autobalance aggregate show-unbalanced-volume-state` command displays information about a volume that is considered by the Auto Balance Aggregate feature.

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>|local}] - Node Name (privilege: advanced)

If this parameter is specified, the display will be limited to only those volumes with a node that matches the specified value.

[-DSID <integer>] - DSID of the Last Volume Queried (privilege: advanced)

If this parameter is specified, the display will be limited to only those volumes with a DSID that matches the specified value.

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

If this parameter is specified, the display will be limited to only those volumes with an aggregate name that matches the specified value.

[-volume-name <text>] - Name of the Volume (privilege: advanced)

If this parameter is specified, the display will be limited to only that volume with a name that matches the specified value.

[-last-threshold-crossed-time <MM/DD/YYYY HH:MM:SS>] - Last Time Threshold Crossed (privilege: advanced)

If this parameter is specified, the display will be limited to only those volumes with a threshold crossing time that matches the specified value.

[-last-placed-time <MM/DD/YYYY HH:MM:SS>] - Last Time Volume Was Moved (privilege: advanced)

If this parameter is specified, the display will be limited to only those volumes with a last time they have been moved that matches the specified value.

[-is-moving {true|false}] - Is Volume Currently Moving (privilege: advanced)

If this parameter is specified as true, the display will be limited to only the moving volumes. If the parameter is specified as false, the display will be limited to only the non-moving volumes.

[-is-quiesced {true|false}] - Is Volume Quiesced (privilege: advanced)

If this parameter is specified as true, the display will be limited to only the quiesced volumes. If the parameter is specified as false, the display will be limited to only the non-quiesced volumes.

[-total-footprint {<integer>[KB|MB|GB|TB|PB]}] - Total Size of the Volume (privilege: advanced)

If this parameter is specified, the display will be limited to only those volumes with a total footprint that matches the specified value.

[-attributes <text>,…] - Volume's Attributes (privilege: advanced)

If this parameter is specified, the display will be limited to only those volumes with attributes that matches the specified value.

[-last-checked <MM/DD/YYYY HH:MM:SS>] - Last Time Volume State Was Checked (privilege: advanced)

If this parameter is specified, the display will be limited to only those volumes with a last time their state was checked that matches the specified value.

Examples

The following example display information about all of the unbalanced volumes that the Auto Balance Aggregate feature is aware of.

```
cluster1::*> autobalance aggregate show-unbalanced-volume-state
              Last Checked On: 3/13/2014 14:32:01
Volume: ro10
                    Footprint: 20.20MB
Last Time Over IOPS Threshold: 3/12/2014 16:20:18
                  Last Placed: 3/11/2014 10:16:04
                   Attributes: Over IOPS Threshold
                               Stabilizing
Volume: test
                    Footprint: 20.20MB
Last Time Over IOPS Threshold: 3/12/2014 16:20:18
                  Last Placed: 3/11/2014 10:16:42
                   Attributes: Over IOPS Threshold
                               In Mirror
                               Stabilizing
```

The following example displays all of the information that the Auto Balance Aggregate feature has collected for all of the unbalanced volumes it is aware of.

```
cluster1::*> autobalance aggregate show-unbalanced-volume-state -instance
                         Node Name: cluster-1-01
   DSID of the Last Volume Queried: 1025
                         Aggregate: aggr_1
                Name of the Volume: ro10
       Last Time Threshold Crossed: 3/12/2014 16:20:18
        Last Time Volume Was Moved: 3/11/2014 10:16:04
        Is Volume Currently Moving: false
                Is Volume Quiesced: false
          Total Size of the Volume: 20.20MB
               Volume's Attributes: Over IOPS Threshold
                                    Stabilizing
Last Time Volume State Was Checked: 3/13/2014 08:20:18
Node Name: cluster-1-01
   DSID of the Last Volume Queried: 1026
                         Aggregate: aggr_1
                Name of the Volume: test
       Last Time Threshold Crossed: 3/12/2014 16:20:18
        Last Time Volume Was Moved: 3/11/2014 10:16:42
        Is Volume Currently Moving: false
                Is Volume Quiesced: false
          Total Size of the Volume: 20.20MB
               Volume's Attributes: Over IOPS Threshold
                                    In Mirror
                                    Stabilizing
Last Time Volume State Was Checked: 3/13/2014 08:20:18
```

autobalance aggregate config modify

Modify the Auto Balance Aggregate feature configuration

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

Description

The `autobalance aggregate config modify` command allows the user to customize the parameters that determine when volumes should be considered for automatic move or recommendation by the Auto Balance Aggregate feature.

Parameters

`[-is-enabled {true|false}]` - Is the Auto Balance Aggregate Feature Enabled (privilege: advanced)

This specifies whether the Auto Balance Aggregate feature is enabled and running.

[-aggregate-unbalanced-threshold-percent <integer>] - Threshold When Aggregate Is Considered Unbalanced (%) (privilege: advanced)

This specifies the space used threshold percentage that will cause the Auto Balance Aggregate feature to consider an aggregate as unbalanced.

[-aggregate-available-threshold-percent <integer>] - Threshold When Aggregate Is Considered Balanced (%) (privilege: advanced)

This specifies the threshold percentage which will determine if an aggregate is a target destination for a move. The Auto Balance Aggregate feature will attempt to move volumes from an unbalanced aggregate until it is under this percentage.

Examples

The following example displays a modification for the default configuration of the Auto Balance Aggregate feature

```
cluster1::*> autobalance aggregate config show
                  Is the Auto Balance Aggregate Feature Enabled: false
          Threshold When Aggregate Is Considered Unbalanced (%): 70
            Threshold When Aggregate Is Considered Balanced (%): 40
cluster1::*> autobalance aggregate config modify -is-enabled true
cluster1::*> autobalance aggregate config show
                  Is the Auto Balance Aggregate Feature Enabled: true
          Threshold When Aggregate Is Considered Unbalanced (%): 70
            Threshold When Aggregate Is Considered Balanced (%): 40
```

autobalance aggregate config show

Display the Auto Balance Aggregate feature configuration

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

Description

The `autobalance aggregate config show` command displays information about parameters that determine when volumes should be considered for automatic move or recommendation by the Auto Balance Aggregate feature.

Examples

The following example displays the default configuration for the Auto Balance Aggregate feature

```
cluster1::*> autobalance aggregate config show
               Is the Auto Balance Aggregate Feature Enabled: false
      Threshold When Aggregate Is Considered Unbalanced (%): 70
        Threshold When Aggregate Is Considered Balanced (%): 40
```

cluster commands

cluster add-node-status-clear-failed

Remove failed nodes from the status list

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

Description

The `cluster add-node-status-clear-failed` command removes failed entries from the status list displayed by the cluster add-node-status command.

Examples

The following example removes failed entries:

```
clus1::> cluster add-node-status
        Node Name        Cluster IP          Status              Error
Reason
        ---------------- ------------------- -------------------
--------------------------
        -                2.2.2.2             failure             The
cluster port MTU
                                                                 values do
not match.
                                                                 Retry
cluster create or
                                                                 join
after updating the
                                                                 cluster
port MTUs to the
                                                                 same
value.
        node1            1.1.1.1             success
        2 entries were displayed.
clus1::> cluster add-node-status-clear-failed
clus1::> cluster add-node-status
        Node Name        Cluster IP          Status              Error
Reason
        ---------------- ------------------- -------------------
--------------------------
        node1            1.1.1.1.            success
```

Related Links

- cluster add-node-status

cluster add-node-status

Show cluster expansion progress

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

Description

The `cluster add-node-status` command displays the progress of the node joining a cluster initiated by using the cluster create command or the cluster add-node command

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-uuid <UUID>] - Node UUID

Select the node that match the specified node UUID.

[-node-name <text>] - Node Name

Select the nodes that match the specified node name.

[-cluster-ip <IP Address>] - IP Address of a Cluster Interface of Node

Select the nodes that match the specified cluster IP.

[-status <Cluster Operation Status>] - The Status of Current Task Being Performed

Select the nodes that match the specified status. This status shows whether the operation is ongoing or complete with success or failure. The various phases that a node goes through are node-initialization, joining-cluster, service-startup, post-cluster-setup and success or failure.

[-failure-msg <text>] - Error Reason

Select the nodes that match the specified error string.

[-last-updated <MM/DD/YYYY HH:MM:SS>] - Last Updated

The date/time stamp of the last update to the status.

Examples

The following example shows the progress of a node add operation:

```
clus1::> cluster add-node-status
        Node Name       Node IP            Status          Error Reason
        --------------- ----------------- --------------
----------------------
        node1           1.1.1.1            success         -
```

Related Links

- cluster create
- cluster add-node

cluster add-node

Expand the cluster by discovering and adding new nodes

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

Description

The `cluster add-node` command discovers and adds new nodes to the cluster. When the `-node-count` parameter is specified, the command attempts to add that many nodes to the cluster. The `-node-ip` parameter can be specified to directly add a node. The `-cluster-ips` parameter can be specified to directly add one or more nodes in parallel. Only one of the `-node-count` , `-node-ip` and `-cluster-ips` parameters can be provided. The system node show-discovered command displays all the nodes discovered on the local network.

> The `node-count` parameter is deprecated and may be removed in a future release of ONTAP. Use the `-cluster-ips` parameter instead.

> The `node-ip` parameter is deprecated and may be removed in a future release of ONTAP. Use the `-cluster-ips` parameter instead.

Parameters

{ `-cluster-ips <IP Address>,…` - List of Cluster Interface IP Addresses of the Nodes Being Added

This parameter contains a comma separated list of cluster interface IP addresses of the nodes in the cluster you are creating. All the nodes specified in the list will be added to the cluster.

| `-retry <true>` - Retry a failed cluster add-node operation

Use this parameter to retry the most recently failed `cluster add-node` command with the originally specified parameters. Retry is not supported if the `cluster add-node` command was originally run with either the *`-node-count`* or *`-node-ip`* parameters.

| `-node-count <integer>` - (DEPRECATED)-Number of Nodes Being Added

Number of nodes to be added to the cluster. If fewer nodes are discovered, all the discovered nodes are added to the cluster and the command will fail since there are fewer nodes than specified. If more nodes are found than the number specified, the command will fail because there is no way to determine which

nodes you intend to add to the cluster.

> The `-node-count` parameter is supported on non-shared architecture platforms only.

| -node-ip <IP Address> - (DEPRECATED)-Cluster IP Address of Node }

Cluster IP address of the node to add. When this parameter is provided, the command directly adds the node.

[-node-names <text>,…] - List of Node Names

This parameter contains a comma separated list of node names of all the nodes in the cluster you are creating. The node names must have an one to one correspondence with *-cluster-ips* parameter. The names provided will be used to rename the nodes once they are added to the cluster.

[-foreground {true|false}] - Foreground Process

When set to *false* the command runs in the background as a job. The default is *true* , which causes the command to return after the operation completes.

[-allow-mixed-version-join <true>] - Allow a Node At a Different Version to Join Cluster

This parameter allows nodes with different, but compatible versions of ONTAP to be added to the cluster. A ONTAP best practice is to add nodes to the cluster that are of the same ONTAP version as the nodes in the cluster, but that may not always be possible.

Examples

The following example adds a node using `-cluster-ips` :

```
cluster1::> cluster add-node -cluster-ips 1.1.1.1, 2.2.2.2
      Use the 'cluster add-node-status' command to see the progress of the
add-node operation.
```

The following example adds 3 nodes using `-node-count` .

```
cluster1::> cluster add-node -node-count 3
      [Job 22] Job succeeded.
```

Related Links

- system node show-discovered

cluster create

Create a cluster

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

Description

The `cluster create` command creates a cluster with one or more nodes. When the `-node-count` parameter is specified, the command attempts to add that many nodes to the cluster. The `-cluster-ips` parameter can be specified to add one or more nodes in parallel. Only one of the `-node-count` and `-cluster-ips` parameters can be provided.

Note that single-node clusters do not require configuring the cluster network. A cluster network interface must be configured before other nodes can join the cluster.

> The `node-count` parameter is deprecated and may be removed in a future release of ONTAP. Use the `-cluster-ips` parameter instead.

Parameters

`[-license <License Code V2>]` - (DEPRECATED)-Base License

> This parameter is deprecated and may be removed in a future release of ONTAP.

Use this optional parameter to specify the base license for the cluster. Obtain this value from your sales or support representative.

`-clustername <text>` - Cluster Name

Use this parameter to specify the name of the cluster you are creating.

- The name must contain only the following characters: A-Z, a-z, 0-9, "-" or "_".
- The first character must be one of the following characters: A-Z or a-z.
- The last character must be one of the following characters: A-Z, a-z or 0-9.
- The maximum supported length is 44 characters.
- The system reserves the following names: "all", "cluster", "local" and "localhost".

| `[-cluster-ips <IP Address>,…]` - List of Cluster Interface IP Addresses of the Nodes Being Added

This parameter contains a comma separated list of cluster interface IP addresses of the nodes in the cluster you are creating. All the nodes specified in the list will be added to the cluster.

| `[-node-count <integer>]` - (DEPRECATED)-Node Count

Use this parameter to specify the number of nodes in the cluster you are creating.

- `-node-count` parameter is supported on non-shared architecture platforms only.

`[-node-names <text>,…]` - List of Node Names

This parameter contains a comma separated list of node names of all the nodes in the cluster you are creating. The node names must have an one to one correspondence with *`-cluster-ips`* parameter. The names provided will be used to rename the nodes once they are added to the cluster.

| `-retry <true>` - Retry a failed cluster create operation }

Use this parameter to retry the most recently failed `cluster create` command with the originally specified parameters. Retry is not supported if the `cluster create` command was originally run with either the *`-node-count`* or *`-node-ip`* parameters.

Examples

The following example creates a cluster named cluster1

```
cluster1::> cluster create -clustername cluster1
```

The following example creates a cluster named cluster1 with node-count 4 on a non-shared architecture platform.

```
cluster1::> cluster create -clustername cluster1 -node-count 4
```

cluster join

(DEPRECATED)-Join an existing cluster using the specified member's IP address or by cluster name

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

Description

> This command is deprecated and may be removed in a future release of ONTAP. Use cluster add-node from a node in the cluster instead.

The `cluster join` command adds a node to an existing cluster. Use the cluster create command to create a cluster if one does not already exist.

Note that a cluster network interface must be configured for the cluster before other nodes can join the cluster.

Parameters

{ -clusteripaddr <IP Address> - IP Address of a Cluster Interface from a Node in the Cluster

Use this parameter to specify the IP address of a cluster interface. This must be the IP address of a cluster interface of a node that is already in the cluster. This parameter is mutually exclusive with the `-cluster -name` parameter.

| -cluster-name <text> - (DEPRECATED)-Cluster Name of the Cluster to Join }

Deprecated. Use this parameter to specify the name of an existing cluster to join.

[-allow-mixed-version-join <true>] - Allow a Node at a Different Version to Join Cluster

This parameter allows nodes with different, but compatible versions of ONTAP to join the cluster. A ONTAP best practice is to join nodes to the cluster that are of the same ONTAP version as the nodes in the cluster, but that may not always be possible.

[-node-name <text>] - Name to Use for the Node in the Cluster

This parameter specifies the name that the node will have when we join it to the cluster.

`-cluster-admin-user <text>` - Cluster Admin User on the Cluster to Join

Use this parameter to specify the username of a user with the "admin" role, for example the user "admin".

Examples

The following example joins the local node to a cluster. The IP address 192.0.2.66 is the address of a cluster interface of a node that already belongs to the cluster.

```
node::> cluster join -clusteripaddr 192.0.2.66
```

Related Links

- cluster add-node
- cluster create

cluster modify

Modify cluster node membership attributes

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

Description

The `cluster modify` command modifies the cluster attributes of a node, including its eligibility to participate in the cluster.

At the advanced privilege level, you can use the command to specify whether a node holds epsilon. Epsilon is an extra fractional vote that enables quorum to form using slightly weaker requirements. For example, two out of four eligible nodes are sufficient to form quorum if one of those two nodes holds epsilon.

Parameters

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

Use this parameter to specify the name of the node to modify. If you do not specify a node, the command runs on the local node.

`[-epsilon {true|false}]` - Epsilon (privilege: advanced)

Use this parameter with the value `true` to specify that the node holds Epsilon in the cluster. Use this parameter with the value `false` to specify that the node does not hold Epsilon in the cluster. In a cluster, only one node can be designated as Epsilon at any given time. You can designate a node as Epsilon to add weight to its voting in a cluster with an even number of nodes.

`[-eligibility {true|false}]` - Eligibility (privilege: advanced)

Use this parameter with the value `true` to specify that the node is eligible to participate in the cluster. Use this parameter with the value `false` to specify that the node is not eligible to participate in the cluster.

If you modify a node as ineligible to participate in the cluster, the command prompts you for confirmation before it runs.

[-skip-quorum-check-before-eligible <true>] - Skip Quorum Check Before Setting Node Eligible (privilege: advanced)

If this parameter is specified, quorum checks will be skipped prior to setting a node eligible. When setting a node to eligible, the operation will continue even if there is a possible data outage due to a quorum issue.

[-skip-quorum-check-before-ineligible <true>] - Skip Quorum Check Before Setting Node Ineligible (privilege: advanced)

If this parameter is specified, quorum checks will be skipped prior to setting a node ineligible. When setting a node to ineligible, the operation will continue even if there is a possible data outage due to a quorum issue.

Examples

This example modifies a node to make it eligible to participate in the cluster.

```
cluster1::*> cluster modify -node node3 -eligibility true
```

The following example removes epsilon from the node named node0 and adds it to the node named node1:

```
cluster1::*> cluster modify -node node0 -epsilon false
cluster1::*> cluster modify -node node1 -epsilon true
```

cluster ping-cluster

Ping remote cluster interfaces and perform RPC server check

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

Description

The `cluster ping-cluster` command probes network connectivity to remote cluster interfaces, and performs an RPC server check.

Parameters

-node <nodename> - Node (privilege: advanced)

Use this parameter to send the ping from the node you specify.

[-use-sitelist {true|false}] - Use Sitelist for Cluster Interfaces (privilege: advanced)

Use this parameter with the value `true` to specify that the command use the sitelist to determine any incomplete cluster IP information. Use this parameter with the value `false` to specify that the command not use the sitelist.

[-skip-rpccheck {true|false}] - Skip RPC Server Check (privilege: advanced)

Use this parameter with the value `true` to specify that the command not perform the rpcinfo check of remote hosts. Use this parameter with the value `false` to specify that the command perform the rpcinfo check. The rpcinfo check checks the status of the RPC servers on the remote hosts. By default, the rpcinfo

check runs on the program number of the portmapper. Use the *-rpc-prognum* parameter to override this default.

[-rpc-prognum <integer>] - RPC Server to Check (privilege: advanced)

Use this parameter to override default behavior and run the rpcinfo check on the program number you specify. By default, the rpcinfo check runs on the program number of the portmapper.

Examples

The following example shows typical output for this command.

```
cluster1::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Local = 10.254.231.102  10.254.91.42
Remote = 10.254.42.25    10.254.16.228
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
................
Detected 1500 byte MTU on 4 path(s):
    Local 10.254.231.102 to Remote 10.254.16.228
    Local 10.254.231.102 to Remote 10.254.42.25
    Local 10.254.91.42 to Remote 10.254.16.228
    Local 10.254.91.42 to Remote 10.254.42.25
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

cluster remove-node

Remove a node from the cluster

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

Description

The `cluster remove-node` command removes a node from a cluster.

Before you can remove a node from a cluster, you must shut down all of the node's shared resources, such as virtual interfaces to clients. If any of a node's shared resources are still active, the command fails. The failure message will display which active resources must be shut down before the node can be removed from the cluster.

Parameters

{ -node <nodename> - Node to Unjoin (privilege: advanced)

Use this parameter to specify the name of the node to remove from the cluster.

| -cluster-ip <IP Address> - IP Address of a Cluster Interface of Node to Unjoin (privilege: advanced) }

Use this parameter to specify the cluster IP of the node to remove from the cluster.

[-skip-quorum-check-before-unjoin <true>] - Skip Quorum Check before Unjoin (privilege: advanced)

If this parameter is specified, quorum checks will be skipped prior to the remove-node command. The operation will continue even if there is a possible data outage due to a quorum issue.

[-skip-last-low-version-node-check <true>] - Skip the Check That Prevents Unjoining the Last Low Versioned Node (privilege: advanced)

This parameter allows the node with lowest version of ONTAP to be removed from the cluster.

Examples

The following example shows how to remove the node named `node4` from the cluster.

```
cluster1::*> cluster remove-node -node node4
```

cluster setup

Setup wizard

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

Description

> Use of this command to join a node to an existing cluster is deprecated and might be removed in a future release of ONTAP. From a node in the cluster use the system node show-discovered command and then use the cluster add-node command.

The `cluster setup` command runs the cluster setup wizard, which can be used to either create a cluster or join a node to an existing cluster. When you run the cluster setup wizard, enter the appropriate information at the prompts. You will be asked to provide the following information to create a cluster:

- Node management interface port, IP address, netmask, default gateway
- Cluster name
- Cluster base license key

> This parameter has been deprecated. It may be removed from a future release of ONTAP.

- Feature license keys

- Cluster administrator's password
- Cluster management interface port, IP address, netmask, default gateway
- DNS domain names
- Name server IP addresses
- Location

You will be asked to provide the following information to join a cluster:

- Node management interface port, IP address, netmask, default gateway
- Cluster IP address

The cluster management interface is used for managing the cluster. It provides one IP address to manage the cluster and will fail over to another node, if necessary. This is the preferred IP address for managing the cluster, but you can also manage the cluster by logging in to the node management IP address of a node in the cluster. Since the cluster management interface must be able to fail over, the port role for the interface must be "data" and typically the best choice for an IP address is one on the data network. The node management interface will not fail over, so an IP address on the management network and a port with the role "node management" is the best choice. Alternatively, you can assign an IP address on the data network to the cluster management interface - if that is better in your network topology - but the port must be a data port. The two examples below illustrate the cluster create and cluster join operations, respectively.

Parameters

Examples

An example of using `cluster setup` to create a cluster with IPv4 addresses is shown below.

```
node::> cluster setup
Welcome to the cluster setup wizard.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
     Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
```

```
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0c]:
Enter the node management interface IP address: 192.0.2.66
Enter the node management interface netmask: 255.255.255.192
Enter the node management interface default gateway: 192.0.2.1
The node management interface has been modified to use port e0c with IP
address 192.0.2.66.
Use your web browser to complete cluster setup by accessing
https://192.0.2.66

Otherwise, press Enter to complete cluster setup using the command line
interface:
Do you want to create a new cluster or join an existing cluster? {create,
join}:
create
Do you intend for this node to be used as a single node cluster? {yes, no}
[no]:

Existing cluster interface configuration found:

Port    MTU     IP              Netmask
e0a     9000    169.254.21.189  255.255.0.0
e0b     9000    169.254.29.73   255.255.0.0

Do you want to use this configuration? {yes, no} [yes]:

Enter the cluster administrator's (username "admin") password:

Retype the password:
Step 1 of 5: Create a Cluster
You can type "back", "exit", or "help" at any question.

Enter the cluster name: cluster1

Creating cluster cluster1

Starting cluster support services .

Cluster cluster1 has been created.
Step 2 of 5: Add Feature License Keys
You can type "back", "exit", or "help" at any question.

Enter an additional license key []:
```

```
Step 3 of 5: Set Up a Vserver for Cluster Administration
You can type "back", "exit", or "help" at any question.
Enter the cluster management interface port [e0d]:
Enter the cluster management interface IP address: 192.0.2.60
Enter the cluster management interface netmask: 255.255.255.192
Enter the cluster management interface default gateway [192.0.2.1]:
A cluster management interface on port e0d with IP address 192.0.2.60 has
been created.  You can use this address to connect to and manage the
cluster.

Enter the DNS domain names: data.example.com
Enter the name server IP addresses: 192.0.2.147
DNS lookup for the admin Vserver will use the data.example.com domain.

Step 4 of 5: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.
SFO is licensed.
SFO will be enabled when the partner joins the cluster.
Step 5 of 5: Set Up the Node
You can type "back", "exit", or "help" at any question.

Where is the controller located []: Sunnyvale
Cluster "cluster1" has been created.

To complete cluster setup, you must join each additional node to the
cluster
by running "system node show-discovered" and "cluster add-node" from a
node in the cluster.
To complete system configuration, you can use either OnCommand System
Manager
or the ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster
management IP address (https://192.0.2.60).

To access the command-line interface, connect to the cluster management
IP address (for example, ssh admin@192.0.2.60).
cluster1::>
```

An example of using `cluster setup` to join a cluster with IPv4 addresses is shown below.

```
node::> cluster setup
Welcome to the cluster setup wizard.

You can enter the following commands at any time:
```

```
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
     Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0c]:
Enter the node management interface IP address: 192.0.2.67
Enter the node management interface netmask: 255.255.255.192
Enter the node management interface default gateway: 192.0.2.1
A node management interface on port e0c with IP address 192.0.2.67 has
been created.

Use your web browser to complete cluster setup by accessing
https://192.0.2.67

Otherwise, press Enter to complete cluster setup using the command line
interface:
Do you want to create a new cluster or join an existing cluster? {create,
join}:
join
Existing cluster interface configuration found:

Port    MTU     IP              Netmask
e0a     9000    169.254.31.170  255.255.0.0
e0b     9000    169.254.115.61  255.255.0.0

Do you want to use this configuration? {yes, no} [yes]:
Step 1 of 3: Join an Existing Cluster
You can type "back", "exit", or "help" at any question.
Enter the IP address of an interface on the private cluster network from
```

```
the
cluster you want to join: 169.254.115.8

Joining cluster at address 169.254.115.8

This node has joined the cluster cluster1.
Step 2 of 3: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.
SFO is licensed.
SFO will be enabled when the partner joins the cluster.
Step 3 of 3: Set Up the Node
You can type "back", "exit", or "help" at any question.
This node has been joined to cluster "cluster1".

To complete cluster setup, you must join each additional node to the
cluster
by running "system node show-discovered" and "cluster add-node" from a
node in the cluster.
To complete system configuration, you can use either OnCommand System
Manager
or the ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster
management IP address (https://192.0.2.60).

To access the command-line interface, connect to the cluster management
IP address (for example, ssh admin@192.0.2.60).
cluster1::>
```

An example of using `cluster setup` to create a cluster with IPv6 addresses is shown below.

```
node::> cluster setup
Welcome to the cluster setup wizard.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
     Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp
Technical
```

```
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination
and resolution, should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0c]:
Enter the node management interface IP address:
FD20:8B1E:B255:814E:8E31:DDF2:8224:593D
You are trying to create an IPv6 address. Creating an IPv6 address will
enable
IPv6 on this node.

Type yes to confirm and continue {yes, no}: yes
Enter the node management interface prefix length: 64
Enter the node management interface default gateway:
FD20:8B1E:B255:814E:0000:0000:0000:0001
A node management interface on port e0c with IP address
fd20:8b1e:b255:814e:8e31:ddf2:8224:593d has been created.

Use your web browser to complete cluster setup by accessing
https://[fd20:8b1e:b255:814e:8e31:ddf2:8224:593d]

Otherwise, press Enter to complete cluster setup using the command line
interface:
Do you want to create a new cluster or join an existing cluster? {create,
join}:
create
Do you intend for this node to be used as a single node cluster? {yes, no}
[no]:
Existing cluster interface configuration found:

Port  MTU  IP    Netmask
e0a  9000  172.31.6.130  255.255.192.0
e0b  9000  172.31.6.131  255.255.192.0

Do you want to use this configuration? {yes, no} [yes]:

Enter the cluster administrator's (username "admin") password:

Retype the password:
Step 1 of 5: Create a Cluster
```

```
You can type "back", "exit", or "help" at any question.

Enter the cluster name: cluster1

Creating cluster cluster1

Starting cluster support services ....

Cluster cluster1 has been created.
Step 2 of 5: Add Feature License Keys
You can type "back", "exit", or "help" at any question.

Enter an additional license key []:
Step 3 of 5: Set Up a Vserver for Cluster Administration
You can type "back", "exit", or "help" at any question.
Enter the cluster management interface port [e0d]:
Enter the cluster management interface IP address:
FD20:8B1E:B255:814E:4827:F558:6150:8D5F
Enter the cluster management interface prefix length: 64
Enter the cluster management interface default gateway
[fd20:8b1e:b255:814e::1]:
A cluster management interface on port e0d with IP address
fd20:8b1e:b255:814e:4827:f558:6150:8d5f has been created.  You can use
this address to connect to and manage the cluster.

Enter the DNS domain names: data.example.com
Enter the name server IP addresses:
FD20:8B1E:B255:814E:8F99:721C:5471:FC1E
DNS lookup for the admin Vserver will use the data.example.com domain.

Step 4 of 5: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.
SFO will not be enabled on a non-HA system.
Step 5 of 5: Set Up the Node
You can type "back", "exit", or "help" at any question.

Where is the controller located []:
Cluster "cluster1" has been created.

To complete cluster setup, you must join each additional node to the
cluster
by running "system node show-discovered" and "cluster add-node" from a
node in the cluster.
To complete system configuration, you can use either OnCommand System
Manager
or the ONTAP command-line interface.
```

```
To access OnCommand System Manager, point your web browser to the cluster
management IP address (https://[fd20:8b1e:b255:814e:4827:f558:6150:8d5f]).

To access the command-line interface, connect to the cluster management
IP address (for example, ssh
admin@fd20:8b1e:b255:814e:4827:f558:6150:8d5f).
cluster1::>
```

An example of using `cluster setup` to join a cluster with IPv6 addresses is shown below.

```
node::> cluster setup
Welcome to the cluster setup wizard.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
     Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination
and resolution, should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0c]:
Enter the node management interface IP address:
FD20:8B1E:B255:814E:3A9F:9EBC:BF57:7A89
You are trying to create an IPv6 address. Creating an IPv6 address will
enable
IPv6 on this node.

Type yes to confirm and continue {yes, no}: yes
Enter the node management interface prefix length: 64
Enter the node management interface default gateway:
```

```
FD20:8B1E:B255:814E:0000:0000:0000:0001
A node management interface on port e0c with IP address
fd20:8b1e:b255:814e:3a9f:9ebc:bf57:7a89 has been created.

Use your web browser to complete cluster setup by accessing
https://[fd20:8b1e:b255:814e:3a9f:9ebc:bf57:7a89]

Otherwise, press Enter to complete cluster setup using the command line
interface:
Do you want to create a new cluster or join an existing cluster? {create,
join}:
join
Existing cluster interface configuration found:

Port  MTU  IP    Netmask
e0a  9000  172.31.6.251  255.255.192.0
e0b  9000  172.31.6.252  255.255.192.0

Do you want to use this configuration? {yes, no} [yes]:
Step 1 of 3: Join an Existing Cluster
You can type "back", "exit", or "help" at any question.
Enter the IP address of an interface on the private cluster network from
the
cluster you want to join: 172.31.6.130

Joining cluster at address 172.31.6.130

Starting cluster support services ....

This node has joined the cluster cluster1.
Step 2 of 3: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.
SFO will not be enabled on a non-HA system.
Step 3 of 3: Set Up the Node
You can type "back", "exit", or "help" at any question.
This node has been joined to cluster "cluster1".

To complete cluster setup, you must join each additional node to the
cluster
by running "system node show-discovered" and "cluster add-node" from a
node in the cluster.
To complete system configuration, you can use either OnCommand System
Manager
or the ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster
```

```
management IP address (https://[fd20:8b1e:b255:814e:4827:f558:6150:8d5f]).

To access the command-line interface, connect to the cluster management
IP address (for example, ssh
admin@fd20:8b1e:b255:814e:4827:f558:6150:8d5f).
cluster1::>
```

Related Links

- system node show-discovered
- cluster add-node

cluster show

Display cluster node members

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

Description

The `cluster show` command displays information about the nodes in a cluster.

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>|local}] - Node

Selects the nodes that match this parameter value.

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

Selects the nodes that match this parameter value.

[-epsilon {true|false}] - Epsilon (privilege: advanced)

Selects the nodes that match this parameter value. In a cluster, only one node can be designated as Epsilon at any given time. You can designate a node as Epsilon to add weight to its voting in a cluster with an even number of nodes.

[-eligibility {true|false}] - Eligibility

Selects the nodes that match this parameter value (true means eligible to participate in the cluster).

[-health {true|false}] - Health

Selects the nodes that match this parameter value (true means online).

Examples

The following example displays information about all nodes in the cluster:

```
cluster1::> cluster show
Node                  Health  Eligibility
--------------------- ------- ------------
node0                 true    true
node1                 true    true
node2                 true    true
node3                 true    true
```

The following example displays information about the node named node1:

```
cluster1::> cluster show -node node1
Node: node1
Eligibility: true
     Health: true
```

cluster agent commands

cluster agent connection create

Create Cloud Agent connection

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

Description

Create a websocket connection to an external manager that will be used to monitor this cluster. The external manager will have access to REST APIs and CCMA counters based upon the privilege of the user creating the connection.

Parameters

`-name <text>` - Name (privilege: advanced)

Specify connection name.

`-destination <text>` - Destination URL (privilege: advanced)

Specify destination URL. For secure WebSocket, this should be amqpwss://<host[:port]>.

`[-use-proxy {true|false}]` - Use Proxy (privilege: advanced)

Establish this connection through the HTTP proxy server associated with this connection's IPspace.

`[-subscribe-address <text>]` - Subscribe Address (privilege: advanced)

AMQP address to which ONTAP will subscribe. Defaults to "ontap.agent.manager".

[-publish-address <text>] - Publish Address (privilege: advanced)

AMQP address to which ONTAP will publish. Defaults to "ontap.agent.cluster".

[-certificate-uuid <UUID>] - Authentication Certificate UUID (privilege: advanced)

Specify UUID of an existing certificate to use for authentication to the destination.

[-csr-url <text>] - Authentication CSR URL (privilege: advanced)

URL to send certificate signing request (CSR) that will be used for authenticating with the destination.

[-csr-token <text>] - Authentication CSR Token (privilege: advanced)

Token for authenticating with CSR URL. Not audited.

[-manager-token <text>] - Manager Token (privilege: advanced)

Specify token supplied by manager application to correlate the connection setup back to the workflow that initiated it.

[-ipspace <IPspace>] - IPspace for Connection (privilege: advanced)

The IPspace to use for the connection.

[-idle-timeout <integer>] - Idle Timeout (in seconds) (privilege: advanced)

AMQP idle timeout. Defaults to 0. If set, the local peer will disconnect if it does not receive AMQP frames within the timeout.

[-address-family {unknown|ipv4|ipv6}] - Address Family (privilege: advanced)

Specify address family.

[-auto-delete-error-minutes {<integer>|-}] - Time to Auto-delete when disconnected (in minutes) (privilege: advanced)

Specify time to live in minutes for cloud agent connection in error state. Connection will be deleted if it stays in error state beyond this time.

Examples

```
cluster-1::> cluster agent connection create -name cloud_agent
-destination amqpwss://manager.example.com
```

cluster agent connection delete

Delete Cloud Agent connection

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

Description

Delete the connection to an external manager.

Parameters

`-name <text>` - Name (privilege: advanced)

Specify connection name.

Examples

```
cluster-1::> cluster agent connection delete -name cloud_agent
```

cluster agent connection modify

Modify Cloud Agent connection

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

Description

Modify agent connection properties.

Parameters

`-name <text>` - Name (privilege: advanced)

Specify connection name.

`[-destination <text>]` - Destination URL (privilege: advanced)

Specify destination URL. For secure WebSocket, this should be amqpwss://<host[:port]>.

`[-use-proxy {true|false}]` - Use Proxy (privilege: advanced)

Establish this connection through the HTTP proxy server associated with this connection's IPspace.

`[-subscribe-address <text>]` - Subscribe Address (privilege: advanced)

AMQP address to which ONTAP will subscribe. Defaults to "ontap.agent.manager".

`[-publish-address <text>]` - Publish Address (privilege: advanced)

AMQP address to which ONTAP will publish. Defaults to "ontap.agent.cluster".

`[-certificate-uuid <UUID>]` - Authentication Certificate UUID (privilege: advanced)

Specify UUID of an existing certificate to use for authentication to the destination.

`[-csr-token <text>]` - Authentication CSR Token (privilege: advanced)

Token for authenticating with CSR URL. Not audited.

`[-idle-timeout <integer>]` - Idle Timeout (in seconds) (privilege: advanced)

AMQP idle timeout. Defaults to 0. If set, the local peer will disconnect if it does not receive AMQP frames within the timeout.

`[-address-family {unknown|ipv4|ipv6}]` - Address Family (privilege: advanced)

Specify address family.

[-auto-delete-error-minutes {<integer>|-}] - Time to Auto-delete when disconnected (in minutes) (privilege: advanced)

Specify time to live in minutes for cloud agent connection in error state. Connection will be deleted if it stays in error state beyond this time.

Examples

```
cluster-1::> cluster agent connection modify -name cloud_agent
-destination amqpwss://new-manager.example.com
```

cluster agent connection show

Display Cloud Agent connections

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

Description

This command displays the connections that have been established to help external managers manage the cluster.

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.

[-name <text>] - Name (privilege: advanced)

Selects the connections that match the supplied name.

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

Selects the connections that match the supplied UUID.

[-destination <text>] - Destination URL (privilege: advanced)

Selects the connections that match the supplied destination.

[-state {connecting|connected|error|disabled|partial|static}] - State (privilege: advanced)

Selects the connections that match the supplied state. The state can have one of the following values:

- connecting - The agent is in the process of establishing the connection.
- connected - The connection is up and active.
- error - The connection is down due to a problem. See error-message and error-time for details.
- disabled - The connection has been administratively disabled.

- partial - The connection has been partially established and is not ready to deliver messages.
- static - The agent is of static instance type and does not require an external connection.

[-use-proxy {true|false}] - Use Proxy (privilege: advanced)

Establish this connection through an HTTP proxy server associated with this connection's IPspace

[-subscribe-address <text>] - Subscribe Address (privilege: advanced)

AMQP address to which ONTAP will subscribe. Defaults to "ontap.agent.manager".

[-publish-address <text>] - Publish Address (privilege: advanced)

AMQP address to which ONTAP will publish. Defaults to "ontap.agent.cluster".

[-certificate-uuid <UUID>] - Authentication Certificate UUID (privilege: advanced)

Selects the connections that match the supplied certificate UUID.

[-csr-url <text>] - Authentication CSR URL (privilege: advanced)

URL to send certificate signing request (CSR) that will be used for authenticating with the destination.

[-manager-token <text>] - Manager Token (privilege: advanced)

Selects the connections that match the supplied manager token.

[-ipspace <IPspace>] - IPspace for Connection (privilege: advanced)

The IPspace to use for the connection

[-error-code <integer>] - Last Error Code (privilege: advanced)

Selects the connections that match the supplied error code.

[-error-message <text>] - Last Error Message (privilege: advanced)

Selects the connections that match the supplied error message.

[-error-time <MM/DD/YYYY HH:MM:SS>] - Last Error Time (privilege: advanced)

Selects the connections that match the supplied error time.

[-node <nodename>] - Node Owning the Connection (privilege: advanced)

The node owning the connection

[-msg-recv-total <integer>] - Total Messages Received from Manager (privilege: advanced)

Total messages received from manager

[-msg-recv-manifest <integer>] - Total Manifest Messages Received from Manager (privilege: advanced)

Total manifest messages received from manager

[-msg-recv-baseline-request <integer>] - Total baseline.request Messages Received from Manager (privilege: advanced)

Total baseline.request messages received from manager

[-msg-recv-counters-request <integer>] - Total counters.request Messages Received from Manager (privilege: advanced)

Total counters.request messages received from manager

[-msg-recv-connection-status <integer>] - Total Connection Status Messages Received from Manager (privilege: advanced)

Total connection status messages received from manager

[-msg-recv-connection-modify <integer>] - Total Connection Modify Messages Received from Manager (privilege: advanced)

Total connection modify messages received from manager

[-msg-recv-agent-status <integer>] - Total agent.status.request Messages Received from Manager (privilege: advanced)

Total agent.status.request Messages received from Manager

[-msg-recv-subscription-create <integer>] - Total Subscription Create Messages Received from Manager (privilege: advanced)

Total subscription create messages received from manager

[-msg-recv-subscription-delete <integer>] - Total Subscription Delete Messages Received from Manager (privilege: advanced)

Total subscription delete messages received from manager

[-msg-recv-unknown <integer>] - Total Unknown Messages Received from Manager (privilege: advanced)

Total unknown messages received from manager

[-msg-send-total <integer>] - Total Messages Sent to Manager (privilege: advanced)

Total messages sent to manager

[-msg-send-manifest-request <integer>] - Total manifest.request Messages Sent to Manager (privilege: advanced)

Total manifest.request messages sent to manager

[-msg-send-baseline <integer>] - Total Baseline Messages Sent to Manager (privilege: advanced)

Total baseline messages sent to manager

[-msg-send-counters <integer>] - Total Counters Messages Sent to Manager (privilege: advanced)

Total counters messages sent to manager

[-msg-send-connection-status <integer>] - Total Connection Status Messages Sent to Manager (privilege: advanced)

Total connection status messages sent to manager

[-msg-send-subscription-status <integer>] - Total Subscription Status Messages Sent to Manager (privilege: advanced)

Total subscription status messages received from manager

[-msg-send-errors <integer>] - Total Error Messages Sent to Manager (privilege: advanced)

Total error messages sent to manager

[-msg-delivery-errors <integer>] - Total Errors Encountered Attempting to Send/receive (privilege: advanced)

Total errors encountered attempting to send/receive

[-msg-send-topicrelay <integer>] - Total ONTAP PubSub Topic Messages Sent to Manager (privilege: advanced)

Total ONTAP PubSub Topic messages sent to manager

[-msg-send-agent-status-details <integer>] - Total agent.status Messages Sent to Manager (privilege: advanced)

Total agent.status Messages sent to Manager

[-msg-connection-total <integer>] - Total Connection Attempts (privilege: advanced)

Total connection attempts

[-msg-connection-errors <integer>] - Total Connection Errors (privilege: advanced)

Total connection errors

[-msg-transport-errors <integer>] - Total Transport Errors (privilege: advanced)

Total transport errors

[-application <text>] - Client Application (privilege: advanced)

Selects the connections that match the supplied application.

[-application-url <text>] - Application URL (privilege: advanced)

Selects the connections that match the supplied application URL.

[-idle-timeout <integer>] - Idle Timeout (in seconds) (privilege: advanced)

AMQP idle timeout. Defaults to 0. If set, the local peer will disconnect if it does not receive AMQP frames within the timeout.

[-address-family {unknown|ipv4|ipv6}] - Address Family (privilege: advanced)

Specify address family.

[-auto-delete-error-minutes {<integer>|-}] - Time to Auto-delete when disconnected (in minutes) (privilege: advanced)

Specify time to live in minutes for cloud agent connection in error state. Connection will be deleted if it stays in error state beyond this time.

Examples

```
cluster-1::> cluster agent connection show
Name          Destination                    State      Application    Use
Proxy
------------- ------------------------------ ---------- --------------
---------
cloudinsights amqpws://172.31.50.251         connected  Cloud Insights
false
```

cluster contact-info commands

cluster contact-info modify

Modify contact information for the cluster

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

Description

The `cluster contact-info modify` command modifies contact information for the cluster administrators. If any values contain spaces, you must enclose those values in quotes.

Use the cluster contact-info show command to display contact information for the cluster administrators.

Parameters

`[-primary-name <text>]` - Name of Primary Contact

Use this parameter to specify the name of the primary contact.

`[-primary-phone <text>]` - Phone Number of Primary Contact

Use this parameter to specify the phone number of the primary contact.

`[-primary-alt-phone <text>]` - Alternate Phone Number of Primary Contact

Use this parameter to specify the alternate phone number of the primary contact.

`[-primary-email <text>]` - Email Address or User ID of Primary Contact

Use this parameter to specify the email address of the primary contact.

`[-secondary-name <text>]` - Name of Secondary Contact

Use this parameter to specify the name of the secondary contact.

`[-secondary-phone <text>]` - Phone Number of Secondary Contact

Use this parameter to specify the phone number of the secondary contact.

`[-secondary-alt-phone <text>]` - Alternate Phone Number of Secondary Contact

Use this parameter to specify the alternate phone number of the secondary contact.

`[-secondary-email <text>]` - Email Address or User ID of Secondary Contact

Use this parameter to specify the email address of the secondary contact.

`[-business-name <text>]` - Business Name

Use this parameter to specify the name of the business responsible for this cluster.

`[-address <text>]` - Business Address

Use this parameter to specify the street address of the business responsible for this cluster.

`[-city <text>]` - City Where Business Resides

Use this parameter to specify the name of the city in which the business is located.

[-state <text>] - State Where Business Resides

Use this parameter to specify the name of the state or province in which the business is located.

[-country <Country Code>] - 2-Character Country Code

Use this parameter to specify the 2-character country code of the country in which the business is located.

[-zip-code <text>] - Postal Code Where Business Resides

Use this parameter to specify the postal or ZIP code area in which the business is located.

Examples

The following example changes the name and phone numbers of the secondary contact person for the cluster.

```
cluster1::> cluster contact-info modify -secondary-name "John Doe"
-secondary-phone 123.555.0156 -secondary-alt-phone 123.555.0178
```

The following example changes the mailing address of the business responsible for the cluster.

```
cluster1::> cluster contact-info modify -address "123 Example Avenue"
-city Exampleville -state "New Example" -zip-code 99999 -country US
```

Related Links

- cluster contact-info show

cluster contact-info show

Display contact information for the cluster

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

Description

The `cluster contact-info show` command displays contact information for the cluster administrators.

Examples

The following example shows example output for this command.

```
cluster1::> cluster contact-info show
Name of Primary Contact : Richard Roe
              Phone Number of Primary Contact : 123.555.0123
    Alternate Phone Number of Primary Contact : 123.555.0145
  Email Address or User Id of Primary Contact : roe@example.com
                   Name of Secondary Contact : John Doe
            Phone Number of Secondary Contact : 123.555.0167
  Alternate Phone Number of Secondary Contact : 123.555.0189
Email Address or User Id of Secondary Contact : doe@example.com
                               Business Name : Example Dot Com
                            Business Address : 123 Example Avenue
                  City Where Business Resides : Exampleville
                 State Where Business Resides : New Example
                     2-Character Country Code : US
           Postal Code Where Business Resides : 99999
```

cluster controller-replacement commands

cluster controller-replacement network displaced-interface delete

Delete network interfaces displaced away from this node by controller-replacement

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

Description

Delete references to logical interfaces that were displaced to another node due to a controller-replacement. A LIF that has been displaced to another node has had its home-node modified to another node because no network ports were available in the same broadcast domain on the original home-node. Deleting this entry does not delete the LIF, it only deletes the entry from the displaced-lifs table indicating that the LIF's current home-node is considered restored.

Parameters

`-node <nodename>` - Node

Selects the node from which the LIF was displaced.

`-vserver <vserver>` - Vserver

Selects the vserver on which the LIF resides.

`-lif-name <text>` - Lif Name

Selects the name of the LIF for which to display displaced information.

Examples

The following example deletes displaced LIF information.

```
cluster1::> cluster controller-replacement network displaced-interface
delete -vserver vs0 -lif lif1
```

cluster controller-replacement network displaced-interface restore-home-node

Restore home node for networked interfaces displaced by controller-replacement

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

Description

Restore the original home node of logical interfaces that were displaced to another node due to a controller-replacement. A LIF that has been displaced to another node has had its home-node modified to another node because no network ports were available in the same broadcast domain on the original home-node. Restoring the home-node attempts to find a home-port on the original home node if a suitable port exists.

Parameters

`-node <nodename>` - Node

Selects the node from which the LIF was displaced.

`-vserver <vserver>` - Vserver

Selects the vserver on which the LIF resides.

`-lif-name <text>` - Lif Name

Selects the name of the displaced LIF to be restored.

Examples

The following example restores the home-node of a displaced LIF.

```
cluster1::> cluster controller-replacement network displaced-interface
restore-home-node -vserver vs0 -lif lif1
```

cluster controller-replacement network displaced-interface show

Display network interfaces displaced away from this node by controller-replacement

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

Description

Display logical interfaces that were displaced to another node due to a controller-replacement. A LIF that has been displaced to another node has had its home-node modified to another node because no network ports were available in the same broadcast domain on the original home-node.

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

Selects the node from which the LIF was displaced.

[-vserver <vserver>] - Vserver

Selects the vserver on which the LIF resides.

[-lif-name <text>] - Lif Name

Selects the name of the LIF for which to display displaced information.

[-original-home-node <nodename>] - Original Home Node

The original home-node that was assigned to the LIF prior to controller-replacement.

[-current-home-node <nodename>] - Current Home Node

The current home-node assigned to the LIF after controller-replacement.

Examples

The following example displays the displaced LIF information.

```
cluster1::> cluster controller-replacement network displaced-interface
show
cluster controller-replacement network displaced-interface show)
riginal           Current
server          LIF Name                Home Node          Home Node
--------------- ----------------------- ------------------
------------------
s0              lif1                    node1              node2
1 entry was displayed.
```

cluster controller-replacement network displaced-vlans delete

Remove VLANs displaces by controller-replacement

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

Description

Delete VLAN tags that were displaced due to a controller-replacement. A VLAN tag that has been displaced is a tag that was based on a network port that either no longer exists, or was moved to a new broadcast domain.

Restoring the vlan-tags re-creates them on the specified network port.

Parameters

`-node <nodename>` - Node

Selects the node on which the displaced vlans reside.

`-port <netport>` - Original Base Port

The original base port where the vlans existed prior to controller-replacement.

Examples

The following example deletes the displaced vlan-tag information.

```
cluster1::> cluster controller-replacement network displaced-vlans delete
-node local -port e0c
```

cluster controller-replacement network displaced-vlans restore

Delete VLANs displaced by controller-replacement

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

Description

Restore VLAN tags that were displaced due to a controller-replacement. A VLAN tag that has been displaced is a tag that was based on a network port that either no longer exists, or was moved to a new broadcast domain. Restoring the vlan-tags re-creates them on the specified network port.

Parameters

`-node <nodename>` - Node

Selects the node on which the displaced vlans reside.

`-port <netport>` - Original Base Port

The original base port where the vlans existed prior to controller-replacement.

`-destination-port <netport>` - Destination Port

The destination port where the vlan-tags will be restored.

Examples

The following example restores vlan-tags displaced from port e0c onto port e0d.

```
cluster1::> cluster controller-replacement network displaced-vlans restore
-node node1 -port e0c -destination-port e0d
```

cluster controller-replacement network displaced-vlans show

Display VLANs displaced by controller-replacement

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

Description

Display VLAN tags that were displaced due to a controller-replacement. A VLAN tag that has been displaced is a tag that was based on a network port that either no longer exists, or was moved to a new broadcast domain. Restoring the vlan-tags re-creates them on the specified network port.

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

Selects the node on which the displaced vlans reside.

[-port <netport>] - Original Base Port

The original base port where the vlans existed prior to controller-replacement.

[-vlan-tags <integer>,…] - Displaced VLANs

The vlan-tags that were assigned to the network port prior to controller-replacement.

Examples

The following example displays the displaced vlan-tag information.

```
cluster1::> cluster controller-replacement network displaced-vlans show
cluster controller-replacement network displaced-vlans show)
riginal
ode      Base Port VLANs
-------- ---------
-------------------------------------------------------
ode1     e0c       100,110,120,300,310,320
1 entry was displayed.
```

cluster date commands

cluster date modify

Modify the current date and time for the nodes in the cluster

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

Description

The `cluster date modify` command sets the time zone, date, and time on every node in the cluster.

Parameters

`[-timezone <Area/Location Timezone>]` - Time Zone

This parameter sets the timezone, specified in the Olson format.

`{ [-date {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm]}]` - Date and Time

This parameter sets the date and time, in the format MM/DD/YYYY HH:MM:SS.

`| [-dateandtime <[[[[cc]yy]mm]dd]hhmm[.ss]]>]` - Date and Time

This parameter sets the date and time information, in the format [[[[[cc]yy]mm]dd]hhmm[.ss]]. The argument for setting the date and time is interpreted as follows:

- cc First 2 digits of the year (e.g., 20 for 2011).
- yy Last 2 digits of year (e.g., 10 for 2010).
- mm Numeric month, a number from 01 to 12.
- dd Day, a number from 01 to 31.
- hh Hour, a number from 00 to 23.
- mm Minute, a number from 00 to 59.
- ss Second, a number from 00 to 59.

If the first two digits of the year are omitted, and the last two digits are greater than 68, a date in the 1900s is used. Otherwise, a date in the 2000s is used. If all four digits of the year are omitted, the default is the current year. If the month or day is omitted, the default is the current month or day, respectively. If the seconds are omitted, the default is set to 00. The system automatically handles the time changes for Daylight Saving and Standard time, and for leap seconds and years.

`| [-u, -utcdateandtime <[[[[cc]yy]mm]dd]hhmm[.ss]]>]` - UTC Date and Time }

This parameter sets the date and time information in Coordinated Universal Time (UTC), in the format [[[[[cc]yy]mm]dd]hhmm[.ss]]. -u is an alias for -utcdateandtime. The argument for setting the date and time is interpreted as follows:

- cc First 2 digits of the year (e.g., 20 for 2011).
- yy Last 2 digits of year (e.g., 10 for 2010).
- mm Numeric month, a number from 01 to 12.
- dd Day, a number from 01 to 31.
- hh Hour, a number from 00 to 23.
- mm Minute, a number from 00 to 59.
- ss Second, a number from 00 to 59.

If the first two digits of the year are omitted, and the last two digits are greater than 68, a date in the 1900s is used. Otherwise, a date in the 2000s is used. If all four digits of the year are omitted, the default is the current year. If the month or day is omitted, the default is the current month or day, respectively. If the

seconds are omitted, the default is set to 00. Time changes for Daylight Saving and Standard time, and for leap seconds and years, are handled automatically.

Examples

The following example sets the date and time to January 1 2011, at 1:00 a.m.:

```
cluster1::> cluster date modify -date "01/01/2011 01:00:00"
```

The following example sets the date and time in the UTC format to May 22, 2011, at 09:25:00 a.m.:

```
cluster1::> cluster date modify -u 201105220925.00.
```

cluster date show

Display the current date and time for the nodes in the cluster

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

Description

The `cluster date show` command displays the time zone, date, and time settings for one or more nodes in the cluster. By default, the command displays date and time settings for all nodes in the cluster.

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.

| [-utc]

Displays date and time information in Coordinated Universal Time (UTC).

| [-utcdate]

Displays date and time information in UTC.

| [-instance] }

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

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

Selects the nodes that match this parameter value.

[-timezone <Area/Location Timezone>] - Time Zone

Selects the nodes that match this parameter value (specified in the Olson format).

[-date {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm]}] - Date and Time

Selects the nodes that match this parameter value.

[-utc-date <MM/DD/YYYY HH:MM:SS>] - UTC Date and Time

Selects the nodes that match this parameter value.

[-dateandtime <[[[[[cc]yy]mm]dd]hhmm[.ss]]>] - Date and Time

Selects the nodes that match this parameter value (interpreted as follows):

- cc First 2 digits of the year (e.g., 20 for 2011).
- yy Last 2 digits of year (e.g., 11 for 2011).
- mm Numeric month, a number from 01 to 12.
- dd Day, a number from 01 to 31.
- hh Hour, a number from 00 to 23.
- mm Minute, a number from 00 to 59.
- ss Second, a number from 00 to 59.

[-u, -utcdateandtime <[[[[[cc]yy]mm]dd]hhmm[.ss]]>] - UTC Date and Time

-u is used as an alias for -utcdateandtime. Selects the nodes that match this parameter value (interpreted as follows):

- cc First 2 digits of the year (e.g., 20 for 2011).
- yy Last 2 digits of year (e.g., 11 for 2011).
- mm Numeric month, a number from 01 to 12.
- dd Day, a number from 01 to 31.
- hh Hour, a number from 00 to 23.
- mm Minute, a number from 00 to 59.
- ss Second, a number from 00 to 59.

Examples

The following example displays the date and time settings for all nodes in the cluster:

```
cluster1::> cluster date show
Node      Date                Timezone
--------- ------------------- -----------------
node0     10/06/2011 09:35:15 America/New_York
node1     10/06/2011 09:35:15 America/New_York
node2     10/06/2011 09:35:15 America/New_York
node3     10/06/2011 09:35:15 America/New_York
4 entries were displayed.
```

cluster date zoneinfo load-from-uri

Load timezone zoneinfo data

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

Description

The `cluster date zoneinfo load-from-uri` command loads a new set of timezone zoneinfo data to replace the version installed in the cluster. Releases of ONTAP software contain the timezone data that is current at the time of release. If a change is made to the timezone between ONTAP releases, then an update can be made to the release data. For instance, if a change is made to when daylight saving time is observed for a country then an update to cluster zoneinfo data may be required.

Only zoneinfo files provided by NetApp for use in ONTAP should be used with this command.

To update the zoneinfo database do the following:

- Download the required zoneinfo file from the NetApp support website.
- Place the file on a local web server accessible without password from the cluster.
- Execute the `cluster date zoneinfo load-from-uri` command, passing the Universal Resource Identifier (URI) of the file as parameter.

> The command need only be executed once for the cluster. The data will be distributed to each node of the cluster.

Parameters

`-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - URI of Timezone Zoneinfo Data

URI of the new zoneinfo file.

Examples

The following example loads a new version of the timezone zoneinfo database to the cluster:

```
cluster1::> cluster date zoneinfo load-from-uri
http://www.example.com/ontap_zoneinfo.zip
```

cluster date zoneinfo show

Display cluster timezone zoneinfo information

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

Description

Display information about the current timezone zoneinfo data.

Examples

The following example shows the zoneinfo information for a cluster:

```
cluster1::> cluster date zoneinfo show
Cluster Zoneinfo Version: 2016f
```

cluster ha commands

cluster ha modify

Modify high-availability configuration of cluster management services

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

Description

The `cluster ha modify` command enables or disables cluster high availability in a two-node cluster. Enable high availability when performing some procedures, such as replacing hardware.

> This command is required to enable high availability if the cluster only has two nodes. Do not run this command in a cluster that has three or more nodes.

> Cluster high availability for two-node clusters differs from the storage failover technology used between two nodes for storage high availability.

Parameters

`[-configured {true|false}]` - HA Configured

Use this parameter with the value `true` to enable high availability mode in the cluster. Use this parameter with the value `false` to disable high availability mode in the cluster.

Examples

The following example enables cluster high availability in a cluster.

```
cluster1::> cluster ha modify -configured true
```

cluster ha show

Show high-availability configuration status for the cluster

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

Description

The `cluster ha show` command displays the high-availability status of the cluster. Cluster high-availability mode applies only to two-node clusters.

Examples

The following example displays the high-availability status for a two-node cluster :

```
cluster1::> cluster ha show
High-Availability Configured: true
```

cluster identity commands

cluster identity modify

Modify the cluster's attributes

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

Description

The `cluster identity modify` command changes a cluster's identity information.

Parameters

`[-name <Cluster name>]` - Cluster Name

Use this parameter to specify a new name for the cluster.

- The name must contain only the following characters: A-Z, a-z, 0-9, "-" or "_".
- The first character must be one of the following characters: A-Z or a-z.
- The last character must be one of the following characters: A-Z, a-z or 0-9.
- The maximum supported length is 44 characters.
- The system reserves the following names: "all", "cluster", "local" and "localhost".

`[-location <text>]` - Cluster Location

Use this parameter to specify the physical location of the cluster. For example, "Lab 5".

`[-contact <text>]` - Cluster Contact

Use this parameter to specify contact information for the cluster, such as a name or e-mail address.

Examples

The following example renames the current cluster to *`cluster2`* :

```
cluster1::> cluster identity modify -name cluster2
```

cluster identity show

Display the cluster's attributes including Name, Serial Number, Cluster UUID, Location and Contact

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

Description

The `cluster identity show` command displays the identity information of the cluster.

Examples

The following example displays the cluster's UUID, name, serial number, location and contact information:

```
cluster1::> cluster identity show
Cluster UUID: 1cd8a442-86d1-11e0-ae1c-123478563412
          Cluster Name: cluster1
 Cluster Serial Number: 1-80-123456
      Cluster Location: Lab2
       Cluster Contact: jsmith@example.com

cluster1::>
```

cluster image commands

cluster image cancel-update

Cancel an update

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

Description

The `cluster image cancel-update` command is used to cancel an update that is in either paused-by-user or paused-on-error state. An update can only be canceled if it is in a paused state.

Examples

The following example displays a cancel-update operation:

```
cluster1::> cluster image cancel-update

Warning: The cancel operation can result in a mixed version
         cluster and/or mixed version HA pair. The cancel
         operation can take several minutes to complete.
Do you want to proceed with the cancel operation? {y|n}: y

Info: Canceling update. It may take a few minutes to finish canceling the
update
```

cluster image pause-update

Pause an update

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

Description

The `cluster image pause-update` command is used to pause a currently running update. The update pauses at the next predefined pause point (for example, after validation, download to the boot device, takeover completion, or giveback completion) which might take some time to reach. When the update reaches the pause point, it transitions into the pause-by-user state.

Examples

The following example displays pause-update operation:

```
cluster1::> cluster image pause-update

Info: Pausing update. It may take a few minutes to finish pausing the
update
```

cluster image resume-update

Resume an update

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

Description

The `cluster image resume-update` command is used to resume an update that is currently paused in paused-by-user or paused-on-error state. If the update is not paused then an error is returned.

Parameters

`[-ignore-post-update-checks-failures {true|false}]` - Ignore Post-update-checks Phase Failures (privilege: advanced)

Specifies whether the post update checks phase warnings and/or errors should be ignored. The default value is false.

`[-skip-firmware-update-checks {true|false}]` - Skip Firmware Update Status Checks

Specifies whether firmware update status check will be skipped or not. By default, firmware update status check will not be skipped.

Examples

The following example shows an resume-update operation:

```
cluster1::> cluster image resume-update

Info: Resuming update...
```

cluster image show-update-history

Display the update history

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

Description

The `cluster image show-update-history` command displays the update history for each node. By default, the command displays the following information:

- Status
- Package version
- Start time
- Completion time
- Component ID
- Previous version
- Updated version

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.

[-component-id <text>] - Component ID

Displays updates for the specified component.

[-start-time <MM/DD/YYYY HH:MM:SS>] - Start Time

Displays updates with the specified start time.

[-package-version <text>] - Package Version

Displays updates for the specified package version.

[-status {successful|canceled|back-out|skip-upgrade}] - Status

Displays updates that completed with the specified status.

[-completion-time <MM/DD/YYYY HH:MM:SS>] - Completion Time

Displays updates with the specified completion time.

[-previous-version <text>] - Previous Version

Displays updates with the specified previous version.

[-updated-version <text>] - Updated Version

Displays updates with the specified updated version.

Examples

The following example displays history of automated nondisruptive updates:

```
cluster1::> cluster image show-update-history
Package    Start       Completion                          Previous
Updated
Status     Version    Time        Time        Component ID  Version
Version
---------- ---------  ----------  ----------  ------------  ---------
---------
canceled   8.3        2/11/2014   2/11/2014   ssan-3240-    8.3       8.3
                      12:05:51    12:05:51    55a
successful 8.3        2/11/2014   2/11/2014   ssan-3240-    8.3       8.3
                      14:23:58    15:02:19    55a
successful 8.3        2/13/2014   2/18/2014   ssan-3240-    8.3       8.3
                      16:48:42    09:45:30    55a
successful 8.3        2/18/2014   2/18/2014   ssan-3240-    8.3       8.3
                      10:33:10    11:02:45    55a
canceled   8.3        2/11/2014   2/11/2014   ssan-3240-    8.3       8.3
                      12:05:51    12:05:51    55b
successful 8.3        2/11/2014   2/11/2014   ssan-3240-    8.3       8.3
                      14:23:58    15:54:43    55b
successful 8.3        2/13/2014   2/18/2014   ssan-3240-    8.3       8.3
                      16:48:42    10:05:02    55b
successful 8.3        2/18/2014   2/18/2014   ssan-3240-    8.3       8.3
                      10:33:10    11:22:02    55b
8 entries were displayed.
```

cluster image show-update-log-detail

Display detailed information about nondisruptive update events

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

Description

The `cluster image show-update-log-detail` command displays detailed information about the currently running and previously run nondisruptive update events. By default, the command displays the following information:

- Node
- Transaction ID
- Time stamp
- Destination node
- Task phase
- Task name
- Task status
- Message

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>|local}] - Node (privilege: advanced)

Displays information only for the specified node.

[-task-id <integer>] - Task Id (privilege: advanced)

Displays information only for the specified task ID.

[-posted-time <MM/DD/YYYY HH:MM:SS>] - Posted Time (privilege: advanced)

Displays information that occurred at the specified time.

[-msg-seq-no <integer>] - Message Sequence (privilege: advanced)

Displays information only for the specified message sequence number.

[-current-pid <integer>] - Process ID (privilege: advanced)

Displays information only for the specified process ID.

[-destination <text>] - Task Target node (privilege: advanced)

Displays information only for the specified destination node.

[-ndu-phase {validation|prereq-updates|ontap-updates|package-management|default-phase|post-update-checks}] - Update phase (privilege: advanced)

Displays information only for the specified phase.

[-task-name {initialize|mount-image|restart-hm|get-health|run-scripts|unmount-image|clear-alert|post-restart-hm|cleanup-rd|synch-image|do-download-job|do-failover-job|do-giveback-job|check-progress|complete-validation|invalid-task|default-task|do-postupdate-checks-task}] - Task Name (privilege: advanced)

Displays information only for the specified task name.

[-status {created|ready-to-run|running|completed|failed|pause_req|paused|paused-error|cancel_req|canceled|resume_req|default_status}] - Status Of Task (privilege: advanced)

Displays information only for items with the specified status.

[-message <text>] - Update Log Message (privilege: advanced)

Displays information only for items with the specified message.

[-msg-type <text>] - Type of Message (privilege: advanced)

Displays information only for items with the specified message type.

[-src-info <text>] - Source Information (privilege: advanced)

Displays information only for items for the specified source.

Examples

The following example displays detailed information automated nondisruptive updates:

```
cluster1::*> cluster image show-update-log-detail
           Time     Dest     Task   Task   Task
Node   TID Stamp    Node     Phase  Name   Status Message
------ --- -------- -------- ------ ------ ------
--------------------------
node1  15  3/19/    MUM      ontap- initia ready- Created Task
           2014              update lize   to-run
           13:52:38          s
node1  15  3/19/    MUM      ontap- initia runnin Updated Task Status
           2014              update lize   g
           13:52:38          s
node1  16  3/19/    node1    ontap- do-    ready- Created Task
           2014              update downlo to-run
           13:52:38          s      ad-job
node1  16  3/19/    node1    ontap- do-    runnin Updated Task Status
           2014              update downlo g
           13:52:39          s      ad-job
node1  17  3/19/    node2    ontap- do-    ready- Created Task
           2014              update downlo to-run
           13:52:38          s      ad-job
node2  17  3/19/    node2    ontap- do-    runnin Updated Task Status
           2014              update downlo g
           13:52:38          s      ad-job
6 entries were displayed.
```

cluster image show-update-log

Display the update transaction log

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

Description

The `cluster image show-update-log` command displays detailed information about the currently running, or previously run nondisruptive updates. By default, the command displays the following information:

- Phase
- Transaction
- Transaction ID
- Component ID
- Time stamp
- Status

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.

[-trans-id <integer>] - Transaction ID (privilege: advanced)

Displays information for the step associated with the specified transaction ID.

[-component-id {<nodename>|local}] - Component ID (privilege: advanced)

Displays information for steps associated with the specified component.

[-phase {validation|prereq-updates|ontap-updates|package-management|default-phase|post-update-checks}] - Transaction Phase (privilege: advanced)

Displays information for steps associated with the specified update phase.

[-trans-name {initialize|mount-image|restart-hm|get-health|run-scripts|unmount-image|clear-alert|post-restart-hm|cleanup-rd|synch-image|do-download-job|do-failover-job|do-giveback-job|check-progress|complete-validation|invalid-task|default-task|do-postupdate-checks-task}] - Transaction Name (privilege: advanced)

Displays information for steps associated with the specified transaction.

[-timestamp <MM/DD/YYYY HH:MM:SS>] - Timestamp (privilege: advanced)

Displays information for steps associated with the specified timestamp.

[-status {waiting|started|completed|paused-on-error|paused-by-user|pause-pending|cancel-pending|canceled|failed}] - Status (privilege: advanced)

Displays information for steps matching the specified status.

Examples

The following example displays information about automated nondisruptive update events:

```
cluster1::*> cluster image show-update-log
                            Trans
Phase         Transaction  Id    Component Id       Time Stamp  Status
------------- -----------  ----- -----------------  -----------
---------------
validation    initialize   50    MUM                2/18/2014   completed
                                                    10:32:57

validation    mount-image  51    node1              2/18/2014   completed
                                                    10:32:52

validation    mount-image  52    node2              2/18/2014   completed
                                                    10:32:53

validation    get-health   53    MUM                2/18/2014   completed
                                                    10:32:53

validation    run-scripts  54    node1              2/18/2014   completed
                                                    10:32:53

validation    run-scripts  55    node2              2/18/2014   completed
                                                    10:32:57

validation    unmount-     56    node1              2/18/2014   completed
              image                                 10:32:57

validation    unmount-     57    node2              2/18/2014   completed
              image                                 10:32:57

validation    complete-    58    MUM                2/18/2014   completed
              validation                            10:32:57

package-      cleanup-     66    node1              3/14/2014   completed
management    package                             09:11:51

package-      cleanup-     67    node2              3/14/2014   completed
management    package                             09:11:51

package-      process-     68    node1              3/14/2014   completed
management    package                             09:13:41

package-      synch-image  69    node2              3/14/2014   completed
management                                        09:14:25

13 entries were displayed.
```

cluster image show-update-progress

Display the update progress

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

Description

The `cluster image show-update-progress` command displays information about the current state of an update. By default, the command displays the following information:

- Update phase
- Status
- Estimated Duration
- Elapsed Duration

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.

[-ndu-phase {validation|prereq-updates|ontap-updates|package-management|default-phase|post-update-checks}] - Update Phase

Displays information about the specified update phase.

[-phase-status {in-progress|waiting|paused-by-user|paused-on-error|completed|canceled|failed|pause-pending|cancel-pending}] - Phase Status

Displays information about progress matching the specified phase status.

[-phase-duration <text>] - Phase Duration

Displays information about progress matching the specified phase duration.

[-phase-comments <text>] - Phase Comments

Displays information about progress matching the specified phase comments.

[-elapsed-duration {<seconds>|[<d> days] <hh>:<mm>[:<ss>]}] - Elapsed duration of the phase

Displays information about progress matching the specified elapsed duration.

[-estimated-duration {<seconds>|[<d> days] <hh>:<mm>[:<ss>]}] - Estimated duration of the phase

Displays information about progress matching the specified estimated duration.

[-phase-description <text>] - Phase Description

Displays information about progress matching the specified phase description.

[-subsystem-name <text>] - Subsystem Name

Displays information about progress matching the specified subsystem name.

[-subsystem-status <text>] - Subsystem Status

Displays information about progress matching the specified subsystem status.

[-subsystem-details <text>] - Subsystem Details

Displays information about progress matching the specified subsystem details.

[-subsystem-action <text>] - Subsystem Action

Displays information about progress matching the specified subsystem action.

Examples

The following example shows the cluster image validate progress of two nodes, nodeA and nodeB after executing the cluster image validation command. On the first screen it shows when the pre-update check status is *in-progress* and on the second screen it shows the pre-update checks status changed to *completed*.

```
cluster1::> cluster image show-update-progress
                                             Estimated       Elapsed
pdate Phase         Status                   Duration        Duration
------------------- ----------------- --------------- ---------------
re-update checks    in-progress              00:10:00        00:00:13

etails:

re-update Check     Status            Error-Action
------------------- ----------------- -------------------------------
```

```
cluster1::> cluster image show-update-progress
pdate Phase         Status                   Duration        Duration
------------------- ----------------- --------------- ---------------
re-update checks    completed                00:10:00        00:01:09

etails:

re-update Check     Status        Error-Action
------------------- ----------- --------------------------
oot Menu Status     Warning      Warning: bootarg.init.bootmenu is
                                 enabled on nodes: nodeA,
                                 nodeB. The boot process
                                 of the nodes will be delayed.
                                 Action: Set the bootarg.init.bootmenu
                                 bootarg to false before proceeding
                                 with the upgrade.
```

```
anual checks        Warning      Warning: Manual validation checks
                                 need to be performed. Refer to the
                                 Upgrade Advisor Plan or the "What
                                 should I verify before I upgrade with
                                 or without Upgrade Advisor" section
                                 in the "Upgrade ONTAP" documentation
                                 for the remaining validation checks
                                 that need to be performed before
                                 update. Failing to do so can result
                                 in an update failure or an I/O
                                 disruption.
                                 Action: Refer to the Upgrade Advisor
                                 Plan or the "What should I verify
                                 before I upgrade with or without
                                 Upgrade Advisor" section in the
                                 "Upgrade ONTAP" documentation for the
                                 remaining validation checks that need
                                 to be performed before update.
FS mounts           Warning      Warning: This cluster is serving NFS
                                 clients. If NFS soft mounts are used,
                                 there is a possibility of frequent
                                 NFS timeouts and race conditions that
                                 can lead to data corruption during
                                 the upgrade.
                                 Action: Use NFS hard mounts, if
                                 possible. To list Vservers running
                                 NFS, run the following command:
                                 vserver nfs show
odes to update   list Warning    Warning: List of nodes to be
                                 upgraded: "nodeA,
                                 nodeB."
                                 Action: To upgrade all selected nodes,
                                 regardless of their current version,
                                 set the (privilege: advanced)
                                 "-skip-nodes-at-target-version"
                                 parameter to "false".
 entries were displayed.
```

The following example shows the cluster image validate progress of two nodes, nodeA and nodeB after executing the cluster image validation command along with option show-validation-details. On the first screen it shows when the pre-update check status is *in-progress* . On the second screen it shows the pre-update checks status changed to *completed* .

```
cluster1::> cluster image show-update-progress
                                          Estimated       Elapsed
pdate Phase         Status                   Duration        Duration
------------------- ----------------- --------------- ---------------
re-update checks    in-progress              00:10:00        00:00:13

etails:

re-update Check     Status            Error-Action
------------------- ----------------- -------------------------------
```

```
cluster1::> cluster image show-update-progress

stimated         Elapsed
pdate Phase         Status                   Duration        Duration
------------------- ----------------- --------------- ---------------
re-update checks    completed                00:10:00        00:01:08

etails:

re-update Check     Status            Error-Action
------------------- ----------------- --------------------------------
MPQ Router and      OK                N/A
roker Config
leanup
ggregate online     OK                N/A
tatus and parity
heck
ggregate plex       OK                N/A
esync status check
pplication          OK                N/A
rovisioning Cleanup
utoboot Bootargs    OK                N/A
tatus
ackend              OK                N/A
onfiguration Status
oot Menu Status     Warning           Warning: bootarg.init.bootmenu is
                                       enabled on nodes: nodeA,
                                       nodeB. The boot process
                                       of the nodes will be delayed.
                                       Action: Set the
                                       bootarg.init.bootmenu
                                       bootarg to false before
                                       proceeding with the upgrade.
```

```
roadcast Domain     OK                N/A
vailability and
niqueness for HA
air status
IFS compatibility   OK                N/A
tatus check
LAM quorum online   OK                N/A
tatus check
PU Utilization      OK                N/A
tatus
apacity licenses    OK                N/A
nstall status check
heck For SP/BMC     OK                N/A
onnectivity To
odes
heck LDAP fastbind  OK                N/A
sers using
nsecure connection.
heck for unsecure   OK                N/A
ex algorithm
onfigurations.
heck for unsecure   OK                N/A
ac configurations.
loud keymanager     OK                N/A
onnectivity check
luster health and   OK                N/A
ligibility status
luster quorum       OK                N/A
tatus check
luster/management   OK                N/A
witch check
ompatible New       OK                N/A
mage Check
urrent system       OK                N/A
ersion check if it
s susceptible to
ossible outage
uring NDU
ata ONTAP Version   OK                N/A
nd Previous
pgrade Status
ata aggregates HA   OK                N/A
olicy check
isk status check    OK                N/A
or failed, broken
r non-compatibility
```

```
uplicate Initiator  OK                N/A
heck
ncryption key       OK                N/A
igration status
heck
xternal             OK                N/A
ey-manager with
egacy KMIP client
heck
xternal keymanager  OK                N/A
ey server status
heck
abricpool Object    OK                N/A
tore Availability
igh Availability    OK                N/A
onfiguration
tatus check
nfinite Volume      OK                N/A
vailibility check
IF failover         OK                N/A
apability status
heck
IF health check     OK                N/A
IF load balancing   OK                N/A
tatus check
IFs is on home      OK                N/A
ode status
ogically over       OK                N/A
llocated DP
olumes check
anual checks that   Warning           Warning: Manual validation checks
an be done using                       need to be performed. Refer to the
pgrade ONTAP                           Upgrade Advisor Plan or the "What
ocumentation                           should I verify before I
                                       upgrade with or without Upgrade
                                       Advisor" section in the
                                       "Upgrade ONTAP" documentation
                                     for the remaining validation checks
                                     that need to be performed before
                                     update. Failing to do so can result
                                     in an update failure or an I/O
                                     disruption.
                                     Action: Refer to the Upgrade
                                     Advisor Plan or the "What should
                                     I verify before I upgrade with
                                     or without Upgrade Advisor"
```

```
                                        section in the "Upgrade ONTAP"
                                        documentation for the remaining
                                        validation checks that need
                                        to be performed before update.
etroCluster         OK                   N/A
onfiguration
tatus check for
ompatibility
inimum number of    OK                   N/A
ggregate disks
heck
ulti-Admin upgrade  OK                   N/A
alidation checks
AE Aggregate and    OK                   N/A
VE Volume
ncryption Check
DMP sessions check  OK                   N/A
FS mounts status    Warning              Warning: This cluster is serving
heck                                     NFS clients. If NFS soft mounts
                                         are used, there is a possibility
                                         of frequent NFS timeouts and race
                                         conditions that can lead to data
                                         corruption during the upgrade.
                                         Action: Use NFS hard mounts, if
                                         possible. To list Vservers running
                                         NFS, run the following command:
                                         vserver nfs show
ame Service         OK                   N/A
onfiguration DNS
heck
ame Service         OK                   N/A
onfiguration LDAP
heck
etwork port         OK                   N/A
eachability check
ode to SP/BMC       OK                   N/A
onnectivity check
odes to update list Warning              Warning: List of nodes to be
                                          upgraded: "nodeA,
                                          nodeB."
                                          Action: To upgrade all selected
                                          nodes, regardless of their
                                          current version, set the
                                          (privilege: advanced)
                                          "-skip-nodes-at-target-version"
                                          parameter to "false".
```

```
KM/KMIP enabled     OK                N/A
ystems - Missing
eys check
NTAP API to REST    OK                N/A
ransition warning
NTAP Image          OK                N/A
apability Status
penSSL 3.0.x        OK                N/A
pgrade validation
heck
penssh 7.2 upgrade  OK                N/A
alidation check
latform Health      OK                N/A
onitor check
re-Update           OK                N/A
onfiguration
erification
DB Replica Health   OK                N/A
heck
eplicated database  OK                N/A
chema consistency
heck
unning Jobs Status  OK                N/A
AN LIF association  OK                N/A
tatus check
AN compatibility    OK                N/A
or manual
onfigurability
heck
AN kernel agent     OK                N/A
tatus check
ecure Purge         OK                N/A
peration Check
helves and Sensors  OK                N/A
heck
napLock Version     OK                N/A
heck
napMirror           OK                N/A
ynchronous
elationship status
heck
napMirror           OK                N/A
ompatibility
tatus check
upported platform   OK                N/A
heck
```

```
arget ONTAP         OK                N/A
elease support for
iberBridge 7500N
heck
pgrade Version      OK                N/A
ompatibility Status
erify all bgp       OK                N/A
eer-groups are in
he up state
erify if a cluster  OK                N/A
anagement LIF
xists
erify that e0M is   OK                N/A
ome to no LIFs
ith high speed
ervices.
olume Conversion    OK                N/A
n Progress Check
olume move          OK                N/A
rogress status
heck
olume online        OK                N/A
tatus check
SCSI target portal  OK                N/A
roups status check
verall Status       Warning           Warning
8 entries were displayed.
```

The following example shows the automated nondisruptive update of two nodes, nodeA and nodeB. In this case, nodeA's update is waiting, nodeB's update is in progress. nodeB's giveback operation is in progress.

```
cluster1::> cluster image show-update-progress
Estimated         Elapsed
Update Phase         Status                   Duration        Duration
-------------------- ----------------- --------------- ---------------
Pre-update checks    completed                00:10:00        00:00:02
Data ONTAP updates   in-progress              01:23:00        00:32:07

Details:

Node name            Status            Status Description
-------------------- ----------------- ----------------------------------
nodeA                waiting
nodeB                in-progress       Performing giveback operation.
3 entries were displayed.

cluster1::>
```

The following example shows the automated nondisruptive update of two nodes, nodeA and nodeB. In this case, automated nondisruptive update is paused-on-error in "Data ONTAP updates" phase. nodeA's update is waiting, nodeB's update is failed. "Status Description" show nodeB's error and action.

```
cluster1:> cluster image show-update-progress
Estimated         Elapsed
Update Phase         Status                   Duration        Duration
-------------------- ----------------- --------------- ---------------
Pre-update checks    completed                00:10:00        00:00:02
Data ONTAP updates   paused-on-error           00:49:00        00:05:21

Details:

Node name            Status            Status Description
-------------------- ----------------- ----------------------------------
nodeA                waiting
nodeB                failed            Error: Takeover of node
                                       "nodeB" is not possible.
                                       Action: Use the "storage failover
                                       show" command to view the cause of
                                       the failure.
2 entries were displayed.

Status: Paused - An error occurred in "Data ONTAP updates" phase. The
non-disruptive update cannot continue until the error has been resolved.
Resolve all issues, then use the "cluster image resume-update" command to
resume the update.

cluster1:>
```

The following example shows that the automated nondisruptive update is paused-on-error in "Post-update checks" update phase and "Status Description" shows the error and action.

```
cluster1::> cluster image show-update-progress
Estimated         Elapsed
Update Phase         Status                 Duration        Duration
-------------------- --------------- --------------- ---------------
Data ONTAP updates   completed              02:19:00        00:00:03
Post-update checks   paused-on-error        00:10:00        00:00:02

Details:

Post-update Check    Status         Error-Action
-------------------- -------------- -------------------------------------
Cluster Quorum       Error          Error: Cluster is not in quorum.
Status
                                    Action: Use the (privilege: advanced)
                                    "cluster ring show" command to verify
                                    all replication unit details.
5 entries were displayed.

Status: Paused - An error occurred in "Post-update checks" phase. The
non-disruptive update cannot continue until the error has been resolved.
Resolve all issues, then use the "cluster image resume-update" command
to resume the update.

cluster1::>
```

The following example shows that the automated nondisruptive update is completed on nodeA and nodeB.

```
cluster1::> cluster image show-update-progress
Estimated         Elapsed
Update Phase         Status                   Duration        Duration
-------------------- ----------------- --------------- ---------------
Pre-update checks    completed                00:10:00        00:00:13
Data ONTAP updates   completed                01:23:00        01:15:11
Post-update checks   completed                00:10:00        00:00:02
3 entries were displayed.

Updated nodes: nodeA, nodeB.
cluster1:>
```

The following example shows the automated update of two-node MetroCluster configuration having clusters cluster_A and cluster_B. In this case, cluster_A's update is waiting and cluster_B's update is in progress. cluster_B's switchback operation is in progress.

```
cluster_A::> cluster image show-update-progress
Estimated         Elapsed
Cluster                                Duration        Duration    Status
------------------------------ --------------- ---------------
----------------
cluster_A                              00:00:00        00:00:00    waiting
cluster_B                              00:00:00        00:06:42    in-
progress

Details: Switchback in progress.

Waiting for partner cluster "sti60-vsim-ucs134f_siteB" to be up.

cluster_A::>
```

The following example shows that the automated update is completed on both cluster_A and cluster_B in two-node MetroCluster configuration.

```
cluster_A::> cluster image show-update-progress
Estimated         Elapsed
Cluster                                Duration        Duration    Status
------------------------------ --------------- ---------------
----------------
cluster_A                              00:00:00        00:20:44
completed
cluster_B                              00:00:00        00:10:43
completed

Details: MetroCluster updated successfully.
cluster_A::>
```

cluster image show

Display currently running image information

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

Description

The `cluster image show` command displays information about the version of ONTAP that is running on each node and the date/time when it was installed. By default, the command displays the following information:

- Node name
- Current version
- Installation date and time

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>|local}] - Node

Displays information about the specified node.

[-version <text>] - Current Version

Displays information about the nodes running the specified version.

[-date <MM/DD/YYYY HH:MM:SS>] - Date Installed

Displays information about the nodes with the specified installation date.

Examples

The following example displays information about currently running images on all nodes of the cluster:

```
cluster1::> cluster image show
                     Current                       Installation
Node                 Version                       Date
----------------     -----------------------       ------------
node1                8.3                           -
node2                8.3                           -
2 entries were displayed.
```

cluster image update

Manage an update

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

Description

The `cluster image update` command is used to initiate a ONTAP update. The update is preceded by a validation of the cluster to ensure that any issues that might affect the update are identified. There are two types of updates of a cluster. A rolling update updates ONTAP one HA pair at a time. This type of update is performed for clusters with fewer than eight nodes or when the `-force-rolling` option is specified for clusters with eight or more nodes. A batch update is used for clusters of eight or more nodes, and performs updates of multiple HA pairs at the same time.

There are predefined points in the update when the update can be paused (either by the user or by an error). These pause points occur after validation, after download to the boot device, after takeover has completed, and after giveback has completed.

Parameters

`-version <text>` - Update Version

Specifies the ONTAP version to use to update the cluster.

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

Specifies the nodes that are to be updated. This parameter is not supported for MetroCluster configurations updates or two-stage upgrades. The node upgrade sequence does not follow the node order entered in `cluster image update` command

`[-estimate-only <true>]` - Estimate Only

Creates a report of the steps that occur during the update without performing them precisely.

`[-pause-after {none|takeover-giveback|all}]` - Update Pause

Specifies that the update should pause at each predefined pause points (for example, after validation, after download to the boot device, after takeover, and after giveback) during the update.

`[-ignore-validation-warning {true|false}]` - Ignore Validation

Specifies that the update should proceed even if the validation reports warnings.

`[-skip-confirmation {true|false}]` - Skip Confirmation

Specifies that a validation that does not detect any error issues should not ask the user to confirm the update but simply proceed with the update.

`[-force-rolling <true>]` - Force Rolling Update

This option is used for clusters with eight or more nodes to specify that a rolling update (one HA pair at a time) should be done. This parameter is not supported for single-node cluster and two-node MetroCluster.

`[-stabilize-minutes <integer>]` - Minutes to stabilize

Specifies the number of minutes that the update should wait after a takeover or giveback is completed. This allows time for the clients to recover from the pause in I/O that occurs during takeover and giveback. This parameter is not supported for single-node cluster.

`[-show-validation-details <true>]` - Shows All Validation Details

Specify to display all validation details. Default: do not display all details.

`[-skip-nodes-at-target-version <true>]` - Skip Nodes at Target Version (privilege: advanced)

When this parameter is *`true`* , only the nodes which are not running the target version are upgraded. Nodes which are currently running the target version are skipped. The default value is *`true`* . This parameter cannot be set to true in MetroCluster configurations.

Examples

The following example shows the cluster image update operation. configuration having node nodeA and nodeB.

```
cluster1::> cluster image update -version 9.16.1

Starting validation for this update...
```

```
It can take several minutes to complete validation...

WARNING: There are additional manual upgrade validation checks that must
be performed after these automated validation checks have completed
successfully.
Refer to the Upgrade Advisor Plan or the "What should I verify before I
upgrade with or without Upgrade Advisor" section in the "Upgrade ONTAP"
documentation for the remaining manual validation checks that need to be
performed before update.
Upgrade ONTAP documentation available at: https://docs.netapp.com/us-
en/ontap/upgrade/index.html
The list of checks are available at: https://docs.netapp.com/us-
en/ontap/upgrade/task_what_to_check_before_upgrade.html
Failing to do so can result in an update failure or an I/O disruption.
Please use Interoperability Matrix Tool (IMT
http://mysupport.netapp.com/matrix)
to verify host system supportability configuration information.

Pre-update Check         Status         Error-Action
------------------------ --------      ________________-----------------
Boot Menu Status         Warning      Warning: bootarg.init.bootmenu
                                      is enabled on nodes:
                                      nodeA,
                                      nodeB. The boot
                                      process of the nodes will be
                                      delayed.
                                      Action: Set the
                                      bootarg.init.bootmenu bootarg
                                      to false before proceeding with
                                      the upgrade.
Manual checks            Warning      Warning: Manual validation
                                      need to be performed.
                                      Refer to the Upgrade Advisor
                                      Plan or the "What should I
                                      verify before I upgrade with or
                                      without Upgrade Advisor"
                                      section in the "Upgrade ONTAP"
                                      documentation for the remaining
                                      validation checks that need to
                                      be performed before update.
                                      Failing to do so can result in
                                      an update failure or an I/O
                                      disruption.
                                      Action: Refer to the Upgrade
                                      Advisor Plan or the "What
                                      should I verify before I
```

```
                                        upgrade with or without Upgrade
                                        Advisor" section in the
                                        "Upgrade ONTAP" documentation
                                        for the remaining validation
                                        checks that need to be
                                        performed before update.
NFS mounts                 Warning      Warning: This cluster is
                                        serving NFS clients. If NFS
                                        soft mounts are used, there is
                                        a possibility of frequent NFS
                                        timeouts and race conditions
                                        that can lead to data
                                        corruption during the upgrade.
                                        Action: Use NFS hard mounts, if
                                        possible. To list Vservers
                                        running NFS, run the following
                                        command: vserver nfs show
Nodes to update list       Warning      Warning: List of nodes to be
                                        upgraded: "nodeA,
                                        nodeB."
                                        Action: To upgrade all selected
                                        nodes, regardless of their
                                        current version, set the
                                        (privilege: advanced)
                                        "-skip-nodes-at-target-version"
                                        parameter to "false".
Overall Status             Warning
5 entries were displayed.
Warning: Validation has reported warnings.

Do you want to continue? {y|n}:
Starting update...
Info: Run the "cluster image show-update-progress" command to check update
status.
```

cluster image validate

Validates the cluster's update eligibility

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

Description

The `cluster image validate` command checks for issues within the cluster that might lead to problems during the update.

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.

[-version <text>] - Update Version

Specifies the ONTAP version to use to validate the cluster.

[-rolling <true>] - Rolling Update

Specify this optional parameter on a cluster with eight or more nodes to perform a rolling-update check. The default is to perform a batch-update check.

> This parameter is only supported on a cluster with eight or more nodes, and is not supported for two-node MetroCluster.

[-nodes {<nodename>|local}] - Nodes

Specifies the nodes that are to be validated. This parameter is not supported for MetroCluster configurations and for two-stage upgrades.

[-show-validation-details <true>] - Shows All Validation Details

Specify to display all validation details. Default: do not display all details.

[-skip-nodes-at-target-version <true>] - Skip Nodes at Target Version (privilege: advanced)

Specifies that the validation will skip the nodes which are already running target version. Default: true.

Examples

The following example shows the validate operation:

```
cluster1::> cluster image validate -version 9.16.1

It can take several minutes to complete validation...

WARNING: There are additional manual upgrade validation checks that must
be performed after these automated validation checks have completed
successfully. Refer to the Upgrade Advisor Plan or the "What should I
verify before I upgrade with or without Upgrade Advisor" section in the
"Upgrade ONTAP" documentation for the remaining manual validation checks
that need to be performed before update.
Upgrade ONTAP documentation available at: https://docs.netapp.com/us-
en/ontap/upgrade/index.html
The list of checks are available at: https://docs.netapp.com/us-
en/ontap/upgrade/task_what_to_check_before_upgrade.html
Failing to do so can result in an update failure or an I/O disruption.
Please use Interoperability Matrix Tool
(IMT  http://mysupport.netapp.com/matrix) to verify host system
supportability configuration information.

Validation checks started. Run the "cluster image show-update-progress"
command to check validation status.
```

The following example shows the validate operation with validation details:

```
cluster1::> cluster image validate -version 9.16.1 -show-validation
-details

It can take several minutes to complete validation...

WARNING: There are additional manual upgrade validation checks that must
be performed after these automated validation checks have completed
successfully.
Refer to the Upgrade Advisor Plan or the "What should I verify before
I upgrade with or without Upgrade Advisor" section in the
"Upgrade ONTAP" documentation for the remaining manual validation
checks that need to be performed before update.
Upgrade ONTAP documentation available at: https://docs.netapp.com/us-
en/ontap/upgrade/index.html
The list of checks are available at: https://docs.netapp.com/us-
en/ontap/upgrade/task_what_to_check_before_upgrade.html
Failing to do so can result in an update failure or an I/O disruption.
Please use Interoperability Matrix Tool
(IMT  http://mysupport.netapp.com/matrix) to verify host system
supportability configuration information.

Validation checks started.  Run the "cluster image show-update-progress"
command to check validation status.
```

cluster image package delete

Remove a package from the cluster image package repository

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

Description

The `cluster image package delete` command deletes the specified version of the package from the package repository. The associated information about the package is also deleted from the update database.

Parameters

`-version <text>` - Version To Be Deleted

Specifies the package version that is to be deleted.

Examples

The following example deletes the package with version 8.3:

```
cluster1::> cluster image package delete -version 8.3

Package Delete Operation Completed Successfully
```

cluster image package get

Fetch a package file from a URL into the cluster image package repository

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

Description

The `cluster image package get` command fetches an ONTAP package file specified by the URL into the cluster. The package is stored in the cluster package repository and the information from the package is stored in the update database.

Parameters

`-url <text>` - Package URL

Specifies the URL from which to get the package.

Examples

The following example displays how to get a package from a URL:

```
cluster1::> cluster image package get -url http://example.com/image.tgz
```

cluster image package show-repository

Display information about packages available in the cluster image package repository

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

Description

The `cluster image package show-repository` command displays the package versions that are in the cluster package repository. By default, the command displays the following information:

- Package version

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.

| [-detail]

This parameter specifies that detailed information should be displayed.

| [-instance] }

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

[-download-ver <text>] - Downloaded Version

Displays packages with the specified download version.

[-component-name <text>,…] - Component Name

Displays packages for the specified component.

[-component-version <text>,…] - Component Version

Displays packages with the specified component version.

[-package-build-time <MM/DD/YYYY HH:MM:SS>] - Package Build Time

Displays packages with the specified build time.

Examples

The following example displays the packages in the cluster package repository:

```
cluster1::> cluster image package show-repository
Package Version Package Build Time
--------------- ------------------
8.3             9/12/2014 10:27:33
```

cluster kernel-service commands

cluster kernel-service show

Display cluster service state in the kernel

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

Description

```
The `cluster kernel-service show` command displays the following
information from the master node for each node in the cluster:
```

- Node name
- The quorum status of that node
- The availability status of that node
- The operational status of that node

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.

[-master-node {<nodename>|local}] - Node (privilege: advanced)

The node in the cluster where the information be being reported from. If this parameter is not specified, the command displays information about all nodes in the cluster.

[-cluster-node <text>] - Cluster Node (privilege: advanced)

The node in the cluster that the information listed is regarding. If this parameter is specified, the command displays information only about the nodes with the specified state value.

[-status-quorum {out-of-quorum|in-quorum}] - Quorum Status (privilege: advanced)

The quorum status of the node specified by `-cluster-node` . If this parameter is specified, the command displays information only about the nodes with the specified state value.

[-status-avail {false|true|unknown}] - Availability Status (privilege: advanced)

The availability status of the node specified by `-cluster-node` . If this parameter is specified, the command displays information only about the nodes with the specified state value.

[-status-oper {unknown|operational|not-operational}] - Operational Status (privilege: advanced)

The operational status of the node specified by `-cluster-node` . If this parameter is specified, the command displays information only about the nodes with the specified state value.

Examples

The following example displays information about all nodes in the cluster:

```
cluster1::*> cluster kernel-service show
Master            Cluster           Quorum        Availability
Operational
Node              Node              Status        Status        Status
----------------- ----------------- ------------- -------------
-------------
cluster1-01       cluster1-01       in-quorum     true
operational
                  cluster1-02       in-quorum     true
operational
2 entries were displayed.

cluster1::*> cluster kernel-service show -instance
Master Node: cluster1-01
       Cluster Node: cluster1-01
      Quorum Status: in-quorum
Availability Status: true
 Operational Status: operational
Master Node: cluster1-01
       Cluster Node: cluster1-02
      Quorum Status: in-quorum
Availability Status: true
 Operational Status: operational
2 entries were displayed.
```

cluster kernel-service config modify

Modify cluster service state in the kernel

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

Description

The `cluster kernel-service config modify` used to manage the cluster kernel-service subsystem for a node.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

The node in the cluster where the configuration is being modified.

`[-kcs-enable-takeover {true|false}]` - Initiated Takeover Enabled (privilege: advanced)

This indicates whether the kernel service subsystem for this node will initiate a takeover of any node determined to be *`out of quorum`* if allowed by the HA subsystem.

Examples

```
cluster1::*> cluster kernel-service config modify -node cluster1-01 -kcs
-enable-core false
```

cluster kernel-service config show

Display cluster service state in the kernel

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

Description

The `cluster kernel-service config show` is used to display the configuration of the cluster kernel service subystem for one or more nodes.

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>|local}]` - Node (privilege: advanced)

The node in the cluster where the information be being reported from. If this parameter is not specified, the command displays information about all nodes in the cluster.

`[-clam-master-id <integer>]` - CLAM Master Node ID (privilege: advanced)

The node ID of the master node for the cluster. If this parameter is specified, the command displays information only about the nodes with the specified state value.

`[-kcs-enable-takeover {true|false}]` - Initiated Takeover Enabled (privilege: advanced)

Indicates if the distributed kernel service subystem will initiate a takeover of any node determined to be *`out of quorum`* if the HA subsystem allows it. If this parameter is specified, the command displays information only about the nodes with the specified state value.

`[-kcs-enabled {true|false}]` - KCS Enabled (privilege: advanced)

The state of the kernel service subsystem on the specified node. If this parameter is specified, the command displays information only about the nodes with the specified state value.

`[-quorum-epoch <integer>]` - Quorum Epoch (privilege: advanced)

The number of quorum changes for this node.

Examples

```
cluster1::*> cluster kernel-service config show
Node                        Failover Core    Master            FSM
                  Enabled  Enabled  Enabled Node              State
----------------- -------- -------- ------- -----------------
-------------------------
cluster1-01       true     false    false   cluster1-01       Master:
Waiting for heartbeat timeout
cluster1-02       true     false    false   cluster1-01       Non-Master:
Waiting for backoff timeout
2 entries were displayed.

cluster1::*> cluster kernel-service config show -instance
Node: cluster1-01
                       Master Node: cluster1-01
                     Master Node ID: 1000
                            Enabled: true
        Initiated Takeover Enabled: false
Initiated Core on Takeover Enabled: false
                 Current FSM State: Master: Waiting for heartbeat timeout
                  Running Version: 1
                     Quorum Epoch: 115
                    Voting Status: false
                    CHAAQ Enabled: true
Node: cluster1-02
                       Master Node: cluster1-01
                     Master Node ID: 1000
                            Enabled: true
        Initiated Takeover Enabled: false
Initiated Core on Takeover Enabled: false
                 Current FSM State: Non-Master: Waiting for backoff
timeout
                   Running Version: 1
                      Quorum Epoch: 115
                     Voting Status: false
                     CHAAQ Enabled: true

2 entries were displayed.
```

cluster log-forwarding commands

cluster log-forwarding create

Create a log forwarding destination

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

Description

The `cluster log-forwarding create` command creates log forwarding destinations for remote logging.

Parameters

`-destination <Remote InetAddress>` - Destination Host

Host name or IPv4 or IPv6 address of the server to forward the logs to.

`[-port <integer>]` - Destination Port

The port that the destination server listen on.

`[-protocol {udp-unencrypted|tcp-unencrypted|tcp-encrypted}]` - Log Forwarding Protocol

The protocols are used for sending messages to the destination. The protocols can be one of the following values:

- _ `udp-unencrypted` _ - User Datagram Protocol with no security
- _ `tcp-unencrypted` _ - Transmission Control Protocol with no security
- _ `tcp-encrypted` _ - Transmission Control Protocol with Transport Layer Security (TLS)

`[-ipspace <IPspace>]` - IPspace of Destination

The IPspace of the destination server.

`[-verify-server {true|false}]` - Verify Destination Server Identity

When this parameter is set to `true` , the identity of the log forwarding destination is verified by validating its certificate. The value can be set to `true` only when the *`tcp-encrypted`* value is selected in the protocol field. When this value is `true` the remote server might be validated by OCSP. The OCSP validation for cluster logs is controlled with the security config ocsp enable -app audit_log and security config ocsp disable -app audit_log .

`[-facility <Syslog Facility>]` - Syslog Facility

The Syslog facility to use for the forwarded logs.

`[-force <true>]` - Skip the Connectivity Test

Normally, the `cluster log-forwarding create` command checks that the destination is reachable via an ICMP ping, and fails if it is not reachable. Setting this value to `true` bypasses the ping check so that the destination can be configured when it is unreachable.

`[-message-format {legacy-netapp|rfc-5424}]` - Syslog Message Format

Use this parameter to specify the message format to be used for Syslog messages.

The `message-format` can be one of the following values:

- *`legacy-netapp`* - A variation of the RFC-3164 Syslog format (format: <PRIVAL>TIMESTAMP HOSTNAME: MSG)
- *`rfc-5424`* - Syslog format as per RFC-5424 (format: <PRIVAL>VERSION TIMESTAMP HOSTNAME: MSG)

Refer to the respective RFCs for detailed information regarding the Syslog message formats. + The default message format is *`legacy-netapp`* .

[-timestamp-format-override {no-override|rfc-3164|iso-8601-utc|iso-8601-local-time}] - Syslog Timestamp Format Override

Use this parameter to override the default timestamp format (based on the `message-format` parameter) used for Syslog messages.

The `timestamp-format-override` can be one of the following values:

- *`no-override`* - Timestamp format based on the `message-format` parameter (*`rfc-3164`* if the message format is *`legacy-netapp`* , *`iso-8601-local-time`* if message format is *`rfc-5424`*)
- *`rfc-3164`* - Timestamp format as per RFC-3164 (format: Mmm dd hh:mm:ss)
- *`iso-8601-utc`* - Timestamp format as per ISO-8601 in UTC (format: YYYY-MM-DDThh:mm:ssZ)
- *`iso-8601-local-time`* - Timestamp format as per ISO-8601 in local time (format: YYYY-MM-DDThh:mm:ss+/-hh:mm)

The default value is *`no-override`* . When this parameter is modified, its value persists even when `message-format` is updated. +

[-hostname-format-override {no-override|fqdn|hostname-only}] - Syslog Hostname Format Override

Use this parameter to override the default hostname format (based on the `message-format` parameter) used for Syslog messages.

The `hostname-format-override` can be one of the following values:

- *`no-override`* - Hostname format based on the `message-format` parameter (*`fqdn`* if the message format is *`rfc-5424`* , *`hostname-only`* if message format is *`legacy-netapp`*)
- *`fqdn`* - Fully Qualified Domain Name (e.g., myhost.example.com)
- *`hostname-only`* - Hostname only, without the domain name (e.g., myhost)

The default value is *`no-override`* . When this parameter is modified, its value persists even when `message-format` is updated. +

Examples

This example causes audit logs to be forwarded to a server at address 192.168.0.1, port 514 with USER facility.

```
cluster1::> cluster log-forwarding create -destination 192.168.0.1 -port
514 -facility user
```

Related Links

- security config ocsp enable
- security config ocsp disable

cluster log-forwarding delete

Delete a log forwarding destination

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

Description

The `cluster log-forwarding delete` command deletes log forwarding destinations for remote logging.

Parameters

`-destination <Remote InetAddress>` - Destination Host

Host name or IPv4 or IPv6 address of the server to delete the forwarding entry for.

`-port <integer>` - Destination Port

The port that the destination server listen on.

Examples

This example deletes the forwarding of all logs to the server at address 1.1.1.1, port 514.

```
cluster1::> cluster log-forwarding delete -destination 1.1.1.1 -port 514
```

cluster log-forwarding modify

Modify log forwarding destination settings

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

Description

The `cluster log-forwarding modify` command modifies log forwarding destinations for remote logging.

Parameters

`-destination <Remote InetAddress>` - Destination Host

The host name or IPv4 or IPv6 address of the server to be modified.

`-port <integer>` - Destination Port

The port that the destinations servers listen on.

`[-ipspace <IPspace>]` - IPspace of Destination

The IPspace of the destination server.

`[-verify-server {true|false}]` - Verify Destination Server Identity

When this parameter is set to `true` , the identity of the log forwarding destination is verified by validating its certificate. The value can be set to `true` only when the *`tcp-encrypted`* value is selected in the protocol field. When this value is `true` the remote server might be validated by OCSP. The OCSP validation for

cluster logs is controlled with the security config ocsp enable -app audit_log and security config ocsp disable -app audit_log .

[-facility <Syslog Facility>] - Syslog Facility

The Syslog facility to use for the forwarded logs.

[-message-format {legacy-netapp|rfc-5424}] - Syslog Message Format

Use this parameter to specify a new Syslog message format to replace the current message format.

[-timestamp-format-override {no-override|rfc-3164|iso-8601-utc|iso-8601-local-time}] - Syslog Timestamp Format Override

Use this parameter to override the default Syslog timestamp format (based on the `message-format` parameter).

[-hostname-format-override {no-override|fqdn|hostname-only}] - Syslog Hostname Format Override

Use this parameter to override the default Syslog hostname format (based on the `message-format` parameter).

Examples

This example modifies the facility of audit logs that are forwarded to the destination server at address 192.168.0.1, port 514.

```
cluster1::> cluster log-forwarding modify -destination 192.168.0.1 -port
514 -facility local1
```

Related Links

- security config ocsp enable
- security config ocsp disable

cluster log-forwarding show

Display log forwarding destinations

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

Description

The `cluster log-forwarding show` command displays log forwarding information:

- Destination (IPv4/IPv6/hostname)
- Port number
- List of log classes
- Facility

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.

[-destination <Remote InetAddress>] - Destination Host

If this optional parameter is specified, the command displays information about the forwarding destinations with the specified host name, IPv4 or IPv6 address.

[-port <integer>] - Destination Port

If this optional parameter is specified, the command displays information about the forwarding destinations with the specified ports.

[-protocol {udp-unencrypted|tcp-unencrypted|tcp-encrypted}] - Log Forwarding Protocol

If this optional parameter is specified, the command displays information about the forwarding destinations with the specified protocols.

[-ipspace <IPspace>] - IPspace of Destination

If this optional parameter is specified, the command displays information about the IPspace to which the forwarding destinations belong.

[-verify-server {true|false}] - Verify Destination Server Identity

If this optional parameter is specified, the command displays information about the forwarding destinations with the specified verify-server values.

[-facility <Syslog Facility>] - Syslog Facility

If this optional parameter is specified, the command displays information about the forwarding destinations with the specified facility.

[-message-format {legacy-netapp|rfc-5424}] - Syslog Message Format

Use this optional parameter to display information about the Syslog destination that has the specified Syslog message format.

[-timestamp-format-override {no-override|rfc-3164|iso-8601-utc|iso-8601-local-time}] - Syslog Timestamp Format Override

Use this optional parameter to display information about the Syslog destination that has the specified Syslog timestamp format override.

[-hostname-format-override {no-override|fqdn|hostname-only}] - Syslog Hostname Format Override

Use this optional parameter to display information about the Syslog destination that has the specified Syslog hostname format override.

Examples

```
The following example displays information about the log forwarding
cluster-1::> cluster log-forwarding show
Verify  Syslog
Destination Host         Port   Protocol        Server  Facility
------------------------ ------ --------------- ------  --------
192.168.0.1              514    udp-unencrypted false   user

cluster-1::> cluster log-forwarding show -instance
Destination Host: 192.168.0.1
                   Destination Port: 514
            Log Forwarding Protocol: udp-unencrypted
             IPspace of Destination: Default
 Verify Destination Server Identity: false
                     Syslog Facility: user
              Syslog Message Format: legacy-netapp
   Syslog Timestamp Format Override: no-override
    Syslog Hostname Format Override: no-override
```

cluster peer commands

cluster peer create

Create a new cluster peer relationship

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

Description

The `cluster peer create` command establishes a peer relationship between two clusters. Cluster peering enables independent clusters to coordinate and exchange data.

Before creating a new cluster peer relationship, make sure that both clusters are individually healthy and that there are no other peer relationships between the two clusters that might interfere with the new relationship.

You can create a cluster peer relationship using the IPv4 or IPv6 protocol. You may not use both protocols within a single relationship.

Use the cluster show and cluster peer show commands on each cluster to display health, peering eligibility, and peering information about the two clusters.

Parameters

`[-peer-addrs <Remote InetAddress>,…]` - Remote Intercluster Addresses

Use this parameter to specify the names or IP addresses of the logical interfaces used for intercluster communication. Separate the addresses with commas.

The addresses you provide here are associated with the remote cluster until you modify or delete the relationship, regardless of whether the addresses are valid. Make sure to provide addresses which you

know will remain available on the remote cluster. You can use the hostnames of the remote cluster's intercluster addresses, the IP addresses of the remote cluster's intercluster LIFs or both.

[-username <text>] - Remote User Name

Use this optional parameter to specify a username that runs a reciprocal `cluster peer create` command on the peered cluster. If you choose not to use the reciprocal creation option, by not supplying a username for reciprocal creation, you must run `cluster peer create` again on the remote cluster to complete the peering relationship.

If you specify the username for the remote cluster, you will be prompted to enter the associated remote password. These credentials are not stored, they are used only during creation to authenticate with the remote cluster and to enable the remote cluster to authorize the peering request. The provided username's profile must have access to the console application in the remote cluster.

Use the security login role show and security login show commands on each cluster to find user names and their privilege levels.

[-no-authentication <true>] - Do Not Use Authentication

Use this optional parameter when omitting the `-username` parameter to indicate that you will create an unauthenticated peering relationship.

[-timeout <integer>] - Operation Timeout (seconds) (privilege: advanced)

Use this optional parameter to specify a timeout value for peer communications. Specify the value in seconds. The default timeout value is 60 seconds.

[-address-family {ipv4|ipv6}] - Address Family of Relationship

Use this optional parameter to specify the address family of the cluster peer relationship. The default is based on existing relationships, existing local intercluster LIFs belonging to a particular address-family, and the addresses supplied to the `cluster peer create` command.

[-offer-expiration {MM/DD/YYYY HH:MM:SS | {1..7}days | {1..168}hours | PnDTnHnMnS | PnW}] - Passphrase Match Deadline

Specifying `cluster peer create` normally creates an offer to establish authentication with a cluster that is a potential cluster peer to this cluster. Such offers expire unless they are accepted within some definite time. Use this optional parameter to specify the date and time at which this offer should expire, the time after which the offer will no longer be accepted.

[-rpc-connect-timeout <integer>] - Timeout for RPC Connect (seconds) (privilege: advanced)

Use this optional parameter to specify a timeout value for the RPC connect during peer communications. Specify the value in seconds. The default timeout value is 10 seconds.

[-update-ping-timeout <integer>] - Timeout for Update Pings (seconds) (privilege: advanced)

Use this optional parameter to specify a timeout value for pings while updating remote cluster information. Specify the value in seconds. The default timeout value is 5 seconds. This parameter applies only to cluster peer relationships using the IPv4 protocol.

[-ipspace <IPspace>] - IPspace for the Relationship

Use this optional parameter to specify the IPspace within which the cluster peering relationship is to operate. The default is to use the 'Default' IPspace.

[-local-name <Cluster name>] - Peer Cluster Local Name

Use this optional parameter to specify a unique local name to identify the remote cluster that is being peered. The local name must conform to the same rules as a cluster name. The default value is the remote cluster name.

[-generate-passphrase <true>] - Use System-Generated passphrase

Use this optional parameter alone to create cluster peer offer for the unidentified clusters or use it along with `-peer-addrs` option to automatically generate the passphrase for the cluster peer operation with the peer cluster.

[-initial-allowed-vserver-peers <Vserver Name>,…] - Vservers allowed for auto peering

Use this optional parameter to specify the list of Vservers for which reciprocal Vserver peering with peer cluster should be enabled. Upon the time of successful peering, Vserver peer permission entries will be created for the peer cluster for the specified list of Vservers.

[-encryption-protocol-proposed {none|tls-psk}] - Encryption Protocol To Be Used In Inter-Cluster Communication

Use this optional parameter to specify how this cluster should use encryption in data connections to the other cluster. Specify 'tls-psk' to specify that TLS should be used with a Pre-Shared Key. Specify 'none' to use no encryption. Where authentication is used, the default is 'tls-psk'. Where authentication is not used, the default is 'none'.

[-applications {snapmirror|flexcache}] - Peering Applications

List of peering applications for initially allowed vservers.

Examples

This example creates a peer relationship between cluster1 and cluster2. This reciprocal create executes the create command on both the local cluster and the remote cluster. The cluster peer create command can use the hostnames of cluster2’s intercluster addresses, the IP addresses of cluster2’s intercluster LIFs, or both. Note that the admin user’s password was typed at the prompt, but was not displayed.

```
cluster1::> cluster peer create -peer-addrs cluster2-d2,10.98.234.246
-username admin

Remote Password:

cluster1::> cluster peer show -instance
Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Available
                       Remote Cluster Name: cluster2
                       Active IP Addresses: 10.98.234.246, 10.98.234.243
                     Cluster Serial Number: 1-80-123456
             Address Family of Relationship: ipv4
       Authentication Status Administrative: no-authentication
          Authentication Status Operational: absent
                         Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows coordinated peer creation. The `cluster peer create` command was issued locally on each cluster. This does not require you to provide the username and password for the remote cluster. There is a password prompt, but if you are logged in as the admin user, you may simply press enter.

```
cluster1::> cluster peer create -peer-addrs cluster2-d2, 10.98.234.246 -no
-authentication

Remote Password:
NOTICE: Addition of the local cluster information to the remote cluster
has
failed with the following error: not authorized for that command. You may
need to repeat this command on the remote cluster.

cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  1-80-123456           Available      absent
```

```
cluster2::> cluster peer create -peer-addrs cluster1-d2 -no-authentication

Remote Password:
NOTICE: Addition of the local cluster information to the remote cluster
has
failed with the following error: not authorized for that command. You may
need to repeat this command on the remote cluster.

cluster2::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster1                  1-80-654321           Available      absent
```

This example shows a reciprocal cluster peer create over IPv6 addresses, that establishes a cluster peer relationship with an IPv6 address family.

```
cluster1::> cluster peer create -peer-addrs
FD20:8B1E:B255:C222:6A17:0BBD:E92C:4523 -username admin

Remote Password:

cluster1::> cluster peer show -instance
Address Family of Relationship: ipv6
                          Peer Cluster Name: cluster2
              Remote Intercluster Addresses:
FD20:8B1E:B255:C222:6A17:0BBD:E92C:4523
         Availability of the Remote Cluster: Available
                        Remote Cluster Name: cluster2
                        Active IP Addresses:
FD20:8B1E:B255:C222:6A17:0BBD:E92C:4523
                      Cluster Serial Number: 1-80-123456
             Address Family of Relationship: ipv6
       Authentication Status Administrative: no-authentication
       Authentication Status Operational: absent
                         Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows creation of an authenticated peering relationship. It is an example of using the coordinated method to create a cluster peer relationship. The `cluster peer create` command is issued locally on each cluster. Before executing this pair of commands, a passphrase to be used with the commands is chosen and given at the prompts. The passphrase can be any text; it is prompted for twice on each cluster, and all four copies of the passphrase must agree. The passphrase does not echo on the screen. The

passphrase must be longer than the minimum length as specified by the `cluster peer policy` on both clusters.

```
cluster1::> cluster peer create -peer-addrs cluster2-d2, 10.98.234.246

Enter the passphrase:
Enter the passphrase again:

Notice: Now use the same passphrase in the "cluster peer create" command
in the
        other cluster.

cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  -                     Unavailable    pending

cluster2::> cluster peer create -peer-addrs cluster1-d2

Enter the passphrase:
Enter the passphrase again:

cluster2::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster1                  1-80-654321           Available      ok
```

This example creates a peer relationship between cluster1 and cluster2. This reciprocal create executes the create command on both the local cluster and the remote cluster. The cluster peer create command can use the hostnames of cluster2's intercluster addresses, the IP addresses of cluster2's intercluster LIFs or both. Note that the admin user's password was typed at the prompt, but was not displayed. The `-local-name` parameter is specified to create a local name used to identify the peer cluster in cases where the name of the peer cluster is not unique or not descriptive.

```
cluster1::> create -peer-addrs 10.98.191.193 -username admin -local-name
locallyUniqueName

cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
locallyUniqueName                1-80-000011           Available
absent

cluster1::> cluster peer show -instance
Peer Cluster Name: locallyUniqueName
              Remote Intercluster Addresses: 10.98.191.193
         Availability of the Remote Cluster: Available
                        Remote Cluster Name: cluster2
                        Active IP Addresses: 10.98.191.193
                      Cluster Serial Number: 1-80-000011
             Address Family of Relationship: ipv4
       Authentication Status Administrative: no-authentication
          Authentication Status Operational: absent
                           Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

The following example create a peer relationship between cluster1 and cluster2 using system-generated passphrases:

```
cluster1::> cluster peer create -peer-addrs 10.98.191.193 -generate
-passphrase
Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
            Peer Cluster Name: cluster2
            Initial Allowed Vserver Peers: -
            Expiration Time: 6/7/2017 09:16:10 +5:30
            Intercluster LIF IP: 10.140.106.185
Warning: make a note of the passphrase - it cannot be displayed again.

cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  -                     Unavailable    pending

cluster2::> cluster peer create -peer-addrs 10.140.106.185

Enter the passphrase:

Clusters cluster1 and cluster2 are peered.

cluster2::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster1                  1-80-654321           Available      ok
```

The following example creates a cluster peer offer from cluster1 for an anonymous cluster using system-generated passphrase with offer expiration period of two days and the cluster2 uses the offer from cluster2 with the system-generated passphrase:

```
cluster1::> cluster peer create -generate-passphrase -offer-expiration
2days
Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
            Peer Cluster Name: Clus_7ShR (temporary generated)
            Initial Allowed Vserver Peers: -
            Expiration Time: 6/9/2017 08:16:10 +5:30
            Intercluster LIF IP: 10.140.106.185
Warning: make a note of the passphrase - it cannot be displayed again.

cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
Clus_7ShR                 -                     Unidentified   pending
cluster2::> cluster peer create -peer-addrs 10.140.106.185

Enter the passphrase:

Clusters cluster1 and cluster2 are peered.

cluster2::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster1                  1-80-654321           Available      ok
```

Cluster "cluster1" creates an offer with `initial-allowed-vserver-peers` option set to Vservers "vs1" and "vs2". And the peer cluster "cluster2" uses the offer and creates peer relationship with cluster1, upon the successful peer relationship establishment, Vserver peer permission entries are created for the Vservers "vs1" and "vs2" in cluster "cluster1" for the peer cluster "cluster2". The following example describes the usage of `initial-allowed-vserver-peers` option in the cluster peer creation workflow:

```
cluster1::> cluster peer create -generate-passphrase -initial-allowed
-vserver-peers vs1,vs2
Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
            Peer Cluster Name: Clus_7ShR (temporary generated)
            Initial Allowed Vserver Peers: vs1,vs2
            Expiration Time: 6/7/2017 09:16:10 +5:30
            Intercluster LIF IP: 10.140.106.185
Warning: make a note of the passphrase - it cannot be displayed again.

cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
Clus_7ShR                 -                     Unidentified   pending
cluster2::> cluster peer create -peer-addrs 10.140.106.185

Enter the passphrase:

Clusters cluster1 and cluster2 are peered.

cluster2::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster1                  1-80-654321           Available      ok
cluster1::> vserver peer permission show

Peer Cluster         Vserver               Applications
-------------------  --------------------  --------------------
cluster2
                     vs1                   snapmirror
                     vs2                   snapmirror
2 entries were displayed.
```

Related Links

- cluster show
- cluster peer show
- security login role show
- security login show

cluster peer delete

Delete a cluster peer relationship

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

Description

The `cluster peer delete` command removes a peering relationship. It removes the relationship records, state data, and all associated jobs.

Before removing the relationship, the command verifies that no resources depend on the relationship. For example, if any SnapMirror relationships exist, the command denies the request to delete the peering relationship. You must remove all dependencies for the deletion to succeed. The `cluster peer delete` command removes only the local instance of the peer relationship. An administrator in the peer cluster must use the `cluster peer delete` command there as well to completely remove the relationship.

Parameters

`-cluster <text>` - Peer Cluster Name

Use this parameter to specify the peering relationship to delete by specifying the name of the peered cluster.

Examples

This example shows a failed deletion due to a SnapMirror dependency.

```
cluster2::> cluster peer delete -cluster cluster1
Error: command failed: Unable to delete peer relationship.  Reason: A
        SnapMirror source exists in this cluster
```

cluster peer modify-local-name

Modify the local name for a cluster peer

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

Description

The `cluster peer modify-local-name` command modifies the local name for a remote cluster. The new local name must be unique among all the local names for the remote clusters with which this cluster is peered.

Parameters

`-name <text>` - Cluster Peer Name

Use this parameter to specify the existing local name for a peer cluster.

`-new-name <Cluster name>` - Cluster Peer Local Name

Use this parameter to specify the new local name of the peer cluster. The new local name must conform to the same rules as a cluster name.

Examples

```
cluster2::> cluster peer modify-local-name -name cluster1 -new-name
cluster1A
```

cluster peer modify

Modify cluster peer relationships

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

Description

The `cluster peer modify` command modifies the attributes of a peering relationship. When you modify a peer relationship and specify `-peer-addrs` , all of the remote addresses must respond, must be intercluster addresses, and must belong to the remote cluster that is being modified; or the modification request is denied.

Parameters

`-cluster <text>` - Peer Cluster Name

Use this parameter to specify the peering relationship to modify by specifying the name of the peered cluster.

`[-peer-addrs <Remote InetAddress>,…]` - Remote Intercluster Addresses

Use this parameter to specify the names or IP addresses of the logical interfaces used for intercluster communication. Separate the addresses with commas. The list of addresses you provide replaces the existing list of addresses.

`[-address-family {ipv4|ipv6}]` - Address Family of Relationship

Use this parameter to specify the address family of the names specified with the peer-addrs parameter.

`[-timeout <integer>]` - Operation Timeout (seconds) (privilege: advanced)

Use this parameter to specify a timeout value for peer communications. Specify the value in seconds.

`[-auth-status-admin {no-authentication|revoked|use-authentication}]` - Authentication Status Administrative

Use this parameter to adjust the authentication in use for the peer relationship. The defined values for this field are as follows.

- no-authentication - The cluster peer relationship uses no authentication.
- use-authentication - The cluster peer relationship is to be authenticated. After you use this value, you will be prompted for a passphrase to be used in determining a new authentication key, just as in the authenticated cluster peer create command or you can use the option generate-passphrase to automatically generate the passphrase.
- revoked - The cluster peer relationship is no longer to be trusted. Peering communication with this cluster peer is suspended until the two clusters set their *`auth-status-admin`* attributes either both to *`no-authentication`* or both to *`use-authentication`* .

Changes should be reflected on both clusters involved in a peering relationship.

[-rpc-connect-timeout <integer>] - Timeout for RPC Connect (seconds) (privilege: advanced)

Use this optional parameter to specify a timeout value for the RPC connect during peer communications. Specify the value in seconds.

[-update-ping-timeout <integer>] - Timeout for Update Pings (seconds) (privilege: advanced)

Use this optional parameter to specify a timeout value for pings while updating remote cluster information. Specify the value in seconds. This parameter applies only to cluster peer relationships using the IPv4 protocol.

[-ipspace <IPspace>] - IPspace for the Relationship

Use this optional parameter to specify that cluster peering communication for this remote cluster is to be done using local intercluster LIFs that are on ports in the named IPspace.

[-generate-passphrase <true>] - Use System-Generated passphrase

Use this optional parameter along with `-auth-status-admin` option's `use-authentication` to automatically generate the passphrase which can be used for cluster peer operation.

[-encryption-protocol-proposed {none|tls-psk}] - Encryption For Inter-Cluster Communication

Use this parameter to adjust the encryption of connections in use for the peer relationship. The defined values for this field are as follows.

- tls-psk - Use TLS with a Pre-Shared Key.
- none - Use no encryption.

Examples

This example modifies the peering relationship to use a new IP address in the remote cluster for intercluster communications and revoke authentication.

View existing cluster peer configuration using following command :

```
cluster1::> cluster peer show -instance
Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Available
                        Remote Cluster Name: cluster2
                        Active IP Addresses: 10.98.234.246, 10.98.234.243
                      Cluster Serial Number: 1-80-123456
                       Remote Cluster Nodes: cluster2-01, cluster2-02
                      Remote Cluster Health: true
                    Unreachable Local Nodes: -
            Address Family of Relationship: ipv4
       Authentication Status Administrative: use-authentication
          Authentication Status Operational: ok
                          Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: tls-psk
```

Modify the cluster peer configuration using following command :

```
cluster1::> cluster peer modify -cluster cluster2 -peer-addrs cluster2-
d2,10.98.234.264 -auth-status-admin revoked
Warning: This will discard the authentication key.

Warning: You are removing authentication from the peering relationship
with
  cluster "cluster2". Use the "cluster peer modify" command on
  cluster "cluster2" with the "-auth-status-admin
  no-authentication" parameter to complete authentication removal from
  the peering relationship.

Do you want to continue?{y|n}:y
```

The following example modifies the peering relationship to use authentication with `-generate-passphrase` option.

```
cluster1::> cluster peer modify -cluster cluster2
  -auth-status-admin use-authentication -generate-passphrase

Notice: Use the below system-generated passphrase in the "cluster peer
modify"
command in the other cluster.
Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
        Expiration Time: 6/7/2017 09:16:10 +5:30
        Peer Cluster Name: cluster2

Warning: make a note of the passphrase - it cannot be displayed again.

Until then, the operational authentication state of the relationship
remains as "pending".

cluster1::> cluster peer offer show

Allowed
Peer Cluster Name     Authentication Creation            Expiration
Vserver Peers
--------------------- -------------- -------------------
------------------- --------------
cluster2              ok-and-offer   6/7/2017 08:16:10   6/7/2017 09:16:10
-

Modify cluster peer relationship in cluster2 with use-authentication
option and use the auto-generated passphrase.

cluster2::> cluster peer modify -cluster cluster2 -auth-status-admin use-
authentication

Notice: Use a auto-generated passphrase or choose a passphrase of 8 or
more characters.
        To ensure the authenticity of the peering relationship, use a
phrase or sequence
        of characters that would be hard to guess.

Enter the passphrase:
Confirm the passphrase:
```

Related Links

- cluster peer create

cluster peer ping

Initiate intercluster connectivity test

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

Description

The `cluster peer ping` command displays the status of the network mesh used by the peering relationship. The command checks the network connection to each remote IP address known by the cluster. This includes all intercluster addresses. It is possible for a known address to be not present during the ping. These addresses are not checked, but the absence is temporary.

The most useful parameters for diagnosing problems are `-count` and `-packet-size` . Use the `-count` and `-packet-size` parameters to diagnose problems similarly to how you use them with the standard ping utility.

To display network connection status within a cluster, use the network ping command.

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.

[-originating-node {<nodename>|local}] - Node that Initiates Ping

Use this parameter to send the ping from the node you specify.

[-destination-cluster <Cluster name>] - Cluster to Ping

Use this parameter to specify the peer cluster you want to ping.

[-destination-node <Peer Node Name>] - Node to Ping in Destination Cluster

Use this parameter to specify a specific node in the destination cluster to ping.

[-ip-address <IP Address>] - Active IP Address

Use this parameter to specify the active IP address you want to ping.

[-count <integer>] - Ping Count

Use this parameter to specify the number of requests to be sent to the destination.

[-status {unknown_node|internal_error|unreachable|session_reachable|interface_reachable}] - Status of Ping Operation

Use this parameter to display only ping results that have the status you specify.

[-timeout <integer>] - Ping Timeout in Seconds

Use this parameter to specify a timeout value in seconds for the ping operation.

[-packet-size <integer>] - Size of Packet

Use this parameter to specify the number of data bytes to be sent in the ping packet.

[-ttl <integer>] - Time to Live/ Number of Hops

Use this parameter to specify the maximum number of network hops a packet may make before it is considered a failure.

[-response-time <double>] - Response Time (ms)

Use this parameter to display only nodes that have the response time (in milliseconds) that you specify. This parameter is most useful when specified with a range of values, such as `>500`

[-destination-node-uuid <UUID>] - Node to Ping in Destination Cluster

Use this parameter to specify a specific node's dblade UUID in the destination cluster to ping.

Examples

This example shows a ping of cluster1 and cluster2 from cluster2. All nodes are reachable.

```
cluster2::> cluster peer ping
Node: node1                  Destination Cluster: cluster2
Destination Node IP Address       Count  TTL RTT(ms) Status
---------------- ---------------- ----- ---- -------
-------------------------
node1            10.98.228.230        1  255   0.209 interface_reachable
node2            10.98.228.234        1  255    0.42 interface_reachable
Node: node2                  Destination Cluster: cluster2
Destination Node IP Address       Count  TTL RTT(ms) Status
---------------- ---------------- ----- ---- -------
-------------------------
node1            10.98.228.230        1  255   0.358 interface_reachable
node2            10.98.228.234        1  255    0.17 interface_reachable
Node: node1                  Destination Cluster: cluster1
Destination Node IP Address       Count  TTL RTT(ms) Status
---------------- ---------------- ----- ---- -------
-------------------------
node3            10.98.229.22         1  255   0.336 interface_reachable
node4            10.98.229.29         1  255   0.354 interface_reachable
Node: node2                  Destination Cluster: cluster1
Destination Node IP Address       Count  TTL RTT(ms) Status
---------------- ---------------- ----- ---- -------
-------------------------
node3            10.98.229.22         1  255   0.354 interface_reachable
node4            10.98.229.29         1  255   0.336 interface_reachable
6 entries were displayed.
```

Related Links

- network ping

cluster peer show

Display peer cluster information

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

Description

The `cluster peer show` command displays information about the peering relationships between the current cluster and other clusters. Cluster peering enables independent clusters to coordinate and exchange data.

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.

[-cluster <text>] - Peer Cluster Name

Selects the peered clusters that match this parameter value.

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

Selects the peered clusters that match this parameter value.

[-peer-addrs <Remote InetAddress>,…] - Remote Intercluster Addresses

Selects the peered clusters that match this parameter value (remote-host name or IP address).

[-availability <availability>] - Availability of the Remote Cluster

Selects the peered clusters that match this parameter value. This parameter can have four different values:

- Available - The peer cluster availability status will be *`Available`* only if all the nodes in the local cluster are able to contact all the nodes in the remote cluster.
- Partial - The peer cluster availability status will be *`Partial`* only if some nodes in the local cluster are not able to contact some or all nodes in the peer cluster.
- Unavailable - The peer cluster availability status will be *`Unavailable`* only if all the nodes in the local cluster are not able to contact any node in the peer cluster.
- Pending - The peer cluster availability status will be *`Pending`* while the system is creating in-memory health data.
- Unidentified - The peer cluster availability status will be Unidentified if the cluster peer offer is created for an anonymous cluster and is unused. When the offer is used, then the availability will get changed to any of the above mentioned status.

Note: If one or more nodes in the local cluster are offline or unreachable, then those nodes are not used to determine the availability status for the remote nodes.

[-rcluster <text>] - Remote Cluster Name

Selects the peered clusters that match this parameter value.

[-ip-addrs <Remote InetAddress>,…] - Active IP Addresses

Selects the peered clusters that match this parameter value.

[-serialnumber <Cluster Serial Number>] - Cluster Serial Number

Selects the peered clusters that match this parameter value.

[-remote-cluster-nodes <text>,…] - Remote Cluster Nodes

Selects the peered clusters that match this parameter value.

[-remote-cluster-health {true|false}] - Remote Cluster Health

Selects the peered clusters that match this parameter value.

- true - This means that there is cluster quorum in the peer cluster.
- false - This means that there is no cluster quorum in the peer cluster.

[-unreachable-local-nodes <text>,…] - Unreachable Local Nodes

Selects the peered clusters that match this parameter value.

[-timeout <integer>] - Operation Timeout (seconds) (privilege: advanced)

Selects the peered clusters that match this parameter value.

[-address-family {ipv4|ipv6}] - Address Family of Relationship

Selects the peered clusters that have a relationship established using this protocol.

[-auth-status-admin {no-authentication|revoked|use-authentication}] - Authentication Status Administrative

Selects the peered clusters that match this parameter value, which must be chosen from the following values.

- no-authentication - The cluster peer relationship uses no authentication.
- use-authentication - The cluster peer relationship is authenticated.
- revoked - The cluster peer relationship is revoked until agreement can be reached.

[-auth-status-operational {ok|absent|pending|expired|revoked|declined|refused|ok-and-offer|absent-but-offer|revoked-but-offer|key-mismatch|intent-mismatch|incapable}] - Authentication Status Operational

Selects the peered clusters that match this parameter value, which must be one of the following values.

- ok - The clusters both use authentication and they have agreed on an authentication key.
- absent - The clusters agree not to use authentication.
- pending - This cluster has made an outstanding offer to authenticate with the other cluster, but agreement has not yet been reached.
- expired - This cluster’s offer to authenticate with the other cluster expired before agreement was reached.
- revoked - This cluster has revoked any prior authentication agreement.

- declined - The other cluster has revoked the authentication agreement and is declining to communicate with this cluster.
- refused - The other cluster actively refuses the communication attempts, perhaps because its part of the peering relationship has been deleted.
- ok-and-offer - The clusters agree on an authentication key and are using it. In addition, this cluster has made an outstanding offer to re-authenticate with the other cluster.
- absent-but-offer - The clusters currently agree that neither side requires authentication of the other, but this cluster has made an outstanding offer to authenticate.
- revoked-but-offer - This cluster has revoked any authentication agreement, but it has made an outstanding offer to authenticate.
- intent-mismatch - The two clusters disagree on whether authentication is required.
- key-mismatch - The two clusters both believe that they are authenticated, but one of the shared secrets has become corrupted.
- incapable - The other cluster is no longer running a version of ONTAP that supports authenticated cluster peering.

`[-rpc-connect-timeout <integer>]` - Timeout for RPC Connect (privilege: advanced)

Selects the peered clusters that match this parameter value.

`[-update-ping-timeout <integer>]` - Timeout for Update Pings (privilege: advanced)

Selects the peered clusters that match this parameter value.

`[-last-updated <MM/DD/YYYY HH:MM:SS>]` - Last Update Time

Selects the peered clusters that match this parameter value.

`[-ipspace <IPspace>]` - IPspace for the Relationship

Selects the peered clusters whose relationships are to cross the named local IPspace. The default value is the IPspace name "Default". In relationships created before ONTAP 8.3.1, the initial value is "-" and is not updated to "Default" until an action is taken on a cluster peer relationship, such as creating, modifying, or deleting a relationship.

`[-encryption-protocol-proposed {none|tls-psk}]` - Proposed Setting for Encryption of Inter-Cluster Communication

Selects the peered clusters that match the value of this parameter, that is, whether they are proposing to encrypt their cross-cluster communication or not.

`[-encryption-protocol {none|tls-psk}]` - Encryption Protocol For Inter-Cluster Communication

Selects the peered clusters that match the value of this parameter, that is, whether they are encrypting their cross-cluster communication or not.

`[-psk-algorithm {akep2|jpake}]` - Algorithm By Which the PSK Was Derived

Selects the peered clusters that match the value of this parameter, that is, which of the following algorithms was used to derive the Pre-Shared Key between the clusters.

- akep2 - The key was obtained through the AKEP2-and-PBKDF2 algorithm.
- jpake - The key was obtained through the J-PAKE algorithm.

Examples

This example shows the output of the cluster peer show command when all nodes in the local cluster are able to contact all nodes in the remote peer cluster. Additionally, the peer relationship is authenticated and operating correctly.

```
cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  1-80-123456           Available      ok
```

Detailed information for this scenario is shown below.

```
cluster1::> cluster peer show -instance
Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Available
                       Remote Cluster Name: cluster2
                       Active IP Addresses: 10.98.234.246, 10.98.234.243
                      Cluster Serial Number: 1-80-123456
                       Remote Cluster Nodes: cluster2-01, cluster2-02
                      Remote Cluster Health: true
                    Unreachable Local Nodes: -
            Address Family of Relationship: ipv4
       Authentication Status Administrative: use-authentication
          Authentication Status Operational: ok
                           Last Update Time: 02/05 21:05:41
              IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows the output of the cluster peer show command when some nodes in the local cluster are not able to contact some or all of the nodes in the remote peer cluster.

```
cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  1-80-123456           Available      ok
```

Detailed information for this scenario is shown below.

```
cluster1::> cluster peer show -instance
Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Partial
                        Remote Cluster Name: cluster2
                        Active IP Addresses: 10.98.234.246, 10.98.234.243
                      Cluster Serial Number: 1-80-123456
                       Remote Cluster Nodes: cluster2-01, cluster2-02
                      Remote Cluster Health: false
                    Unreachable Local Nodes: -
             Address Family of Relationship: ipv4
       Authentication Status Administrative: use-authentication
          Authentication Status Operational: ok
                           Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows the output of the cluster peer show command when some nodes in the local cluster cannot be contacted from the node where the command is executed, but all the other nodes including node on which command is executed are able to contact all nodes in the remote peer cluster.

```
cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  1-80-123456           Available      ok
```

Detailed information for this scenario is shown below.

```
cluster1::> cluster peer show -instance
                          Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Available
                        Remote Cluster Name: cluster2
                        Active IP Addresses: 10.98.234.246, 10.98.234.243
                      Cluster Serial Number: 1-80-123456
                       Remote Cluster Nodes: cluster2-01, cluster2-02
                      Remote Cluster Health: true
                    Unreachable Local Nodes: cluster1-01
             Address Family of Relationship: ipv4
       Authentication Status Administrative: use-authentication
          Authentication Status Operational: ok
                           Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows the output of the cluster peer show command when some nodes in the local cluster cannot be contacted from the node where the command is executed, and the node on which command is executed is also not able to contact the remote peer cluster.

```
cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  1-80-123456           Unavailable    ok
```

Detailed information for this scenario is shown below.

```
cluster1::> cluster peer show -instance
Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Unavailable
                        Remote Cluster Name: cluster2
                        Active IP Addresses: 10.98.234.246, 10.98.234.243
                      Cluster Serial Number: 1-80-123456
                       Remote Cluster Nodes: cluster2-01, cluster2-02
                      Remote Cluster Health: -
                    Unreachable Local Nodes: cluster1-01
             Address Family of Relationship: ipv4
       Authentication Status Administrative: use-authentication
          Authentication Status Operational: ok
                           Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows the output of the cluster peer show command when all the nodes in the local cluster are not able to contact any nodes in the remote peer cluster.

```
cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  1-80-123456           Unavailable    ok
```

Detailed information for this scenario is shown below.

```
cluster1::> cluster peer show -instance
Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Unavailable
                        Remote Cluster Name: cluster2
                        Active IP Addresses: 10.98.234.246, 10.98.234.243
                      Cluster Serial Number: 1-80-123456
                       Remote Cluster Nodes: cluster2-01, cluster2-02
                      Remote Cluster Health: -
                    Unreachable Local Nodes: -
             Address Family of Relationship: ipv4
       Authentication Status Administrative: use-authentication
          Authentication Status Operational: ok
                           Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows the output of the `cluster peer show` command while the system is creating in-memory health data.

```
cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
cluster2                  1-80-123456           Pending        ok
```

Detailed information for this scenario is shown below.

```
cluster1::> cluster peer show -instance
Peer Cluster Name: cluster2
              Remote Intercluster Addresses: cluster2-d2, 10.98.234.246
         Availability of the Remote Cluster: Pending
                       Remote Cluster Name: cluster2
                       Active IP Addresses: 10.98.234.246, 10.98.234.243
                     Cluster Serial Number: 1-80-123456
                      Remote Cluster Nodes: -
                     Remote Cluster Health: -
                   Unreachable Local Nodes: -
            Address Family of Relationship: ipv4
      Authentication Status Administrative: use-authentication
         Authentication Status Operational: ok
                         Last Update Time: 02/05 21:05:41
              IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

This example shows the output of the `cluster peer show` command for the offer created for an anonymous cluster:

```
cluster1::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
Authentication
------------------------- --------------------- --------------
--------------
Clus_4gHR                 -                     Unidentified   pending

Detailed information for this scenario is shown below.

cluster1::> cluster peer show -instance
Peer Cluster Name: Clus_4gHR
              Remote Intercluster Addresses: -
         Availability of the Remote Cluster: Unidentified
                        Remote Cluster Name: Clus_4gHR
                        Active IP Addresses: 10.98.234.246, 10.98.234.243
                      Cluster Serial Number: -
                       Remote Cluster Nodes: -
                      Remote Cluster Health: -
                   Unreachable Local Nodes: -
             Address Family of Relationship: ipv4
       Authentication Status Administrative: use-authentication
          Authentication Status Operational: ok
                           Last Update Time: 02/05 21:05:41
               IPspace for the Relationship: Default
 Encryption for Inter-Cluster Communication: none
```

cluster peer connection show

Show current peering connections for a cluster

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

Description

The `cluster peer connection show` command displays information about the current TCP connections and how they are supporting the set of peering relationships.

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.

[-cluster-name <text>] - Remote Cluster

Selects the connections associated with the named peered cluster.

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

Selects the connections hosted by the given node.

[-connection-type {mgmt-client|mgmt-server|data}] - Cluster Peering Connection Type

Selects the connections of the named type. This parameter can have one of three different values:

- Mgmt-client - Management-plane client connections, created so that this node may make management-plane requests of other nodes.
- Mgmt-server - Management-plane server connections, over which this node services requests made by other nodes' mgmt-client connections.
- Data - Connections made between data-planes of different nodes.

[-index <integer>] - Index of Connection

Selects the connections with the given index value.

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

Selects the connections to the cluster with the given cluster UUID.

[-auth-status-admin {no-authentication|revoked|use-authentication}] - Authentication Status Administrative

Selects connections to the peered clusters whose intended authentication matches this parameter value.

[-auth-status-operational {ok|absent|pending|expired|revoked|declined|refused|ok-and-offer|absent-but-offer|revoked-but-offer|key-mismatch|intent-mismatch|incapable}] - Authentication Status Operational

Selects connections to the peered clusters whose authentication state matches this parameter value.

[-authenticated {true|false}] - Authenticated

Selects connections that have been authenticated, or not, according to this parameter value.

[-port <integer>] - Network Port

Selects the connections whose port matches this parameter value.

[-idle <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Idle Time

Selects the connections whose idle times match this parameter value.

[-address <IP Address>] - Remote Network Address

Selects the connections that have this parameter value as the remote network address.

[-psk-algorithm {akep2|jpake}] - PSK Algorithm

Selects the connections for clusters whose PSKs were generated by the named algorithm. This parameter can have one of two different values:

- akep2 - The original algorithm, PBKDF2 followed by AKEP2.
- jpake - Newer algorithm: J-PAKE.

Examples

This example shows the output of the cluster peer connection show command.

```
cluster1::> cluster peer connection show
Cluster Node               Connection Type  Auth  Encrypt Idle  Remote
Address
------- -----------------  ---------------  ----- ------- ----
--------------
cluster2                      Authentication: ok
        node1
                                       data  true  true   6s
10.10.10.100
                                       data  true  true   7s
10.10.10.100
                                       data  true  true   20s
10.10.10.100
                                       data  true  true   11s
10.10.10.100
                                       data  true  true   7s
10.10.10.100
                                       data  true  true   7s
10.10.10.100
                                       data  true  true   11s
10.10.10.200
                                       data  true  true   11s
10.10.10.200
                                       data  true  true   48s
10.10.10.200
                                       data  true  true   48s
10.10.10.200
                                       data  true  true   37s
10.10.10.200
                                       data  true  true   37s
10.10.10.200
12 entries were displayed.
```

cluster peer health show

Check peer cluster health

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

Description

The `cluster peer health show` command displays information about the health of the nodes in peer clusters from the perspective of the nodes in the local cluster. The command obtains health information by

performing connectivity and status probes of each peer cluster's nodes from each node in the local cluster.

To enable quick access to remote cluster health information, remote cluster health status is periodically checked and cached. These cached results enable users and system features to quickly assess the availability of remote resources. By default, this command accesses cached results. Use the *-bypass-cache*true option to force a current, non-cached check of remote cluster health.

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.

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

Selects the node that matches this parameter value.

[-destination-cluster <Cluster name>] - Peer Cluster

Selects the cluster that matches this parameter value.

[-destination-node <Peer Node Name>] - Peer Node

Selects the node that matches this parameter value.

[-destination-cluster-uuid <UUID>] - Peer UUID

Selects the cluster that matches this parameter value.

[-data-ping {unknown_node|internal_error|unreachable|session_reachable|interface_reachable}] - Status of Data Ping Operation

Selects the nodes that match this parameter value.

[-icmp-ping {unknown_node|internal_error|unreachable|session_reachable|interface_reachable}] - Status of ICMP Ping Operation

Selects the nodes that match this parameter value.

[-node-health {true|false}] - RDB Health of the Node

Selects the nodes that match this parameter value (true means healthy).

[-cluster-health {true|false}] - Cluster Health

Selects the nodes that match this parameter value (true means healthy).

[-availability {true|false}] - Communication Indicator

Selects the nodes that match this parameter value (true means communicating).

[-bypass-cache {true|false}] - Bypass Cache and Determine Health

Bypasses cached results to determine current cluster health (true means bypass the cache). Cached results may not be current, but they are displayed more quickly.

[-last-updated <MM/DD/YYYY HH:MM:SS>] - Last Update Time

Selects the nodes that match this parameter value.

Examples

The following example shows typical output for this command in a cluster of two nodes that has a peer cluster of two nodes.

```
cluster1::> cluster peer health show
Node       Cluster-Name                 Node-Name
             Ping-Status               RDB-Health Cluster-Health
Availability
---------- --------------------------- --------- ---------------
------------
node1
           cluster2                    node3
             Data: interface_reachable
             ICMP: -                   true      true            true
                                       node4
             Data: interface_reachable
             ICMP: -                   true      true            true
node2
           cluster2                    node3
             Data: interface_reachable
             ICMP: -                   true      true            true
                                       node4
             Data: interface_reachable
             ICMP: -                   true      true            true
4 entries were displayed.
```

The following example shows detailed health information for node3 in cluster2 from the perspective of node1 in cluster1.

```
cluster1::> cluster peer health show -originating-node node1 -destination
-cluster cluster2 -destination-node node3 -instance
Local Node: node1
                 Peer Cluster: cluster2
                    Peer Node: node3
                    Peer UUID: 5e4befb2-1f36-11d0-98c9-123476563412
Status of Data Ping Operation: interface_reachable
Status of ICMP Ping Operation: -
       RDB health of the node: true
               Cluster Health: true
      Communication Indicator: true
             Last Update Time: 02/06 18:58:38
```

cluster peer offer cancel

Cancel the outstanding offer to authenticate with a peer cluster

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

Description

The `cluster peer offer cancel` command cancels an outstanding offer to authenticate with a potentially peered cluster. After the command completes, the given cluster can no longer establish authentication using the given authentication offer.

Parameters

`-cluster <text>` - Peer Cluster Name

Use this parameter to specify which offer should be cancelled, by specifying the name of the cluster to which the offer is extended.

Examples

The following example cancels the authentication offer to cluster2.

```
cluster1::> cluster peer offer cancel -cluster cluster2
```

cluster peer offer modify

Modify an outstanding offer to authenticate with a peer cluster

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

Description

The `cluster peer offer modify` command modifies the outstanding offer to authenticate with a potentially peered cluster. Every authentication offer has an expiration time, after which the offer will no longer be honored. This command is used to change that expiration time. To cancel the offer outright, use the cluster peer offer cancel command instead.

Parameters

`-cluster <text>` - Peer Cluster Name

Use this parameter to specify the offer to be modified by indicating the name of the cluster to which it has been extended.

`[-offer-expiration {MM/DD/YYYY HH:MM:SS | {1..7}days | {1..168}hours | PnDTnHnMnS | PnW}]` - Authentication Offer Expiration Time

Use this parameter to specify the new expiration time for the offer.

`[-initial-allowed-vserver-peers <Vserver Name>,…]` - Vservers Initially Allowed for Peering

Use this optional parameter to specify the list of Vservers for which reciprocal Vserver peering with peer cluster should be enabled.

Examples

This example modifies the expiration time for the authentication offer to push it out by an hour.

```
cluster1::> cluster peer offer show
Peer Cluster Name       Authentication Creation            Expiration
----------------------- -------------- -------------------
-------------------
cluster2                absent_but_offer
                                       7/23/2013 14:45:47  7/23/2013
15:45:47

cluster1::> cluster peer offer modify -cluster cluster2 -offer-expiration
"7/23/2013 16:45:47"

cluster1::> cluster peer offer show
Peer Cluster Name       Authentication Creation            Expiration
----------------------- -------------- -------------------
-------------------
cluster2                absent_but_offer
                                       7/23/2013 14:45:47  7/23/2013
16:45:47
```

Related Links

- cluster peer offer cancel

cluster peer offer show

Display outstanding offers to authenticate with a peer cluster

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

Description

The `cluster peer offer show` command displays information about authentication offers still pending with potential peer clusters. By default, the command displays information about all unexpired offers made by the local cluster.

To display detailed information about a specific offer, run the command with the `-cluster` parameter.

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.

[-cluster <text>] - Peer Cluster Name

Selects the offer that matches this parameter value.

[-auth-status-operational {ok|absent|pending|expired|revoked|declined|refused|ok-and-offer|absent-but-offer|revoked-but-offer|key-mismatch|intent-mismatch|incapable}] - Authentication Status Operational

Selects the offers that match this parameter value.

[-offer-creation <MM/DD/YYYY HH:MM:SS>] - Authentication Offer Creation Time

Selects the offers that match this parameter value.

[-offer-expiration {MM/DD/YYYY HH:MM:SS | {1..7}days | {1..168}hours | PnDTnHnMnS | PnW}] - Authentication Offer Expiration Time

Selects the offers that match this parameter value.

[-initial-allowed-vserver-peers <Vserver Name>,…] - Vservers Initially Allowed for Peering

Selects the offers that match this parameter value.

[-offer-creator <text>] - Authentication Offer Creator

Selects the offers that match this parameter value.

[-encryption-protocol-proposed {none|tls-psk}] - Encryption Protocol to Be Used For the Relationship

Selects the offers that match this parameter value.

Examples

The following example displays information about the outstanding authentication offers:

```
cluster1::> cluster peer offer show

Allowed
Peer Cluster Name       Authentication Creation            Expiration
Vserver Peers
----------------------- -------------- -------------------
------------------- --------------
cluster2                absent_but_offer 7/11/2013 22:22:52 7/11/2013
23:22:52 vs1,vs2
```

cluster peer policy modify

Modify the policy configuration for the cluster peering service

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

Description

The `cluster peer policy modify` command modifies the prevailing policy settings. One setting governs whether unauthenticated cluster peer relationships can exist. The other setting specifies a minimum length for passphrases.

Parameters

`[-is-unauthenticated-access-permitted {true|false}]` - Is Unauthenticated Cluster Peer Access Permitted

Use this parameter to specify whether unauthenticated peering relationships are allowed to exist. Setting the parameter value to *`true`* allows such relationships to exist. Setting the value to *`false`* prevents both the creation of unauthenticated peering relationships as well as the modification of existing peering relationships to be unauthenticated. Setting the value to *`false`* is not possible if the cluster currently is in any unauthenticated relationships.

`[-passphrase-minlength <integer>]` - Passphrase Length Minimum

Use this parameter to specify a minimum length for passphrases as given to the cluster peer create or cluster peer modify commands in the future. The default value for this parameter is 8.

`[-is-unencrypted-access-permitted {true|false}]` - Is Unencrypted Cluster Peer Access Permitted

Use this parameter to specify whether peering relationships that do not use encryption are allowed to exist. Setting the parameter value to *`true`* allows such relationships to exist. Setting the value to *`false`* prevents the creation of unauthenticated peering relationships and the modification of existing peering relationships to be unauthenticated, as well as preventing unencrypted peering relationships from being created and the modification of existing peering relationships to be unencrypted. Setting the value to *`false`* is not possible if the cluster currently is in any unauthenticated or unencrypted relationships.

Examples

This example modifies the peering policy to disallow unauthenticated intercluster communications.

```
cluster1::> cluster peer policy show
Is Unauthenticated Cluster Peer Communication Permitted:  true
                       Minimum Length for a Passphrase:  8

cluster1::> cluster peer policy modify -is-unauthenticated-access
-permitted false

cluster1::> cluster peer policy show
Is Unauthenticated Cluster Peer Communication Permitted:  false
                       Minimum Length for a Passphrase:  8
```

Related Links

- cluster peer create
- cluster peer modify

cluster peer policy show

Display the policy configuration for the cluster peering service

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

Description

The `cluster peer policy show` command displays the prevailing cluster peer authentication policy. There are two policies at present: one to control whether any cluster peer relationships can be unauthenticated, and one to control the minimum length for a passphrase. If the policy is set to preclude unauthenticated peering relationships, then unauthenticated relationships cannot be created inadvertently. Passphrases of less than the minimum length may not be used. By default, this minimum length is set to 8, so passphrases must be 8 characters long or longer.

Examples

This example shows the cluster peer policy when unauthenticated relationships may not be created inadvertently.

```
cluster1::> cluster peer policy show
Is Unauthenticated Cluster Peer Communication Permitted:  false
                        Minimum Length for a Passphrase:  9
    Is Unencrypted Cluster Peer Communication Permitted: true
```

cluster quorum-service commands

cluster quorum-service options modify

Modify the settings for cluster quorum-service

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

Description

The `cluster quorum-service options modify` command modifies the values of cluster quorum services options.

Parameters

`[-ignore-quorum-warning-confirmations {true|false}]` - Whether or Not Warnings Are Enabled (privilege: advanced)

Specifies whether cluster quorum warnings and confirmations should be ignored when cluster operations could negatively impact cluster quorum:

- Halting a node (system node halt)
- Rebooting a node (system node reboot)
- Issuing a planned takeover (storage failover takeover)

The default setting is false.

Examples

The following example shows the usage of this command:

```
cluster1::> set advanced

Warning: These advanced commands are potentially dangerous; use them only
when
         directed to do so by NetApp personnel.
         Do you want to continue? {y|n}: y

cluster1::*> cluster quorum-service options modify -ignore-quorum-warning
-confirmations true
```

Related Links

- system node halt
- system node reboot
- storage failover takeover

cluster quorum-service options show

Display the settings for cluster quorum-service

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

Description

The `cluster quorum-service options show` command displays the values of cluster quorum services options.

Examples

The following example demonstrates showing the state of ignore-quorum-warning-confirmations when it is false and true.

```
cluster1::*> set advanced

Warning: These advanced commands are potentially dangerous; use them only
when
         directed to do so by NetApp personnel.
         Do you want to continue? {y|n}: y

cluster1::*> cluster quorum-service options show
Ignore Quorum Warning Confirmations
 -----------------------------------
false
```

cluster ring commands

cluster ring show

Display cluster node member's replication rings

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

Description

The `cluster ring show` command displays a cluster's ring-replication status. Support personnel might ask you to run this command to assist with troubleshooting.

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>|local}]` - Node (privilege: advanced)

Selects the rings that match this parameter value.

`[-unitname {mgmt|vldb|vifmgr|bcomd|crs|availd}]` - Unit Name (privilege: advanced)

Selects the rings that match this parameter value. Possible values are:

- `mgmt` - The management application
- `vldb` - The volume location database
- `vifmgr` - The virtual-interface manager
- `bcomd` - The SAN management daemon
- `crs` - The configuration replication service

[-online {master|secondary|offline}] - Status (privilege: advanced)

Selects the rings that match this parameter value.

[-epoch <integer>] - Epoch (privilege: advanced)

Selects the rings that match this parameter value.

[-master <nodename>] - Master Node (privilege: advanced)

Selects the rings that match this parameter value.

[-local <nodename>] - Local Node (privilege: advanced)

Selects the rings that match this parameter value.

[-db-epoch <integer>] - DB Epoch (privilege: advanced)

Selects the rings that match this parameter value.

[-db-trnxs <integer>] - DB Transaction (privilege: advanced)

Selects the rings that match this parameter value.

[-num-online <integer>] - Number Online (privilege: advanced)

Selects the rings that match this parameter value.

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

Selects the rings that match this parameter value.

Examples

The following example displays information about all replication rings in a two-node cluster:

```
cluster1::*> cluster ring show
Node      UnitName Epoch    DB Epoch DB Trnxs Master    Online
--------- -------- -------- -------- -------- --------- ---------
node0     mgmt     1        1        1068     node0     master
node0     vldb     1        1        98       node0     master
node0     vifmgr   1        1        350      node0     master
node0     bcomd    1        1        56       node0     master
node0     crs      1        1        88       node0     master
node1     mgmt     1        1        1068     node0     secondary
node1     vldb     1        1        98       node0     secondary
node1     vifmgr   1        1        350      node0     secondary
node1     bcomd    1        1        56       node0     secondary
node1     crs      1        1        88       node0     secondary
10 entries were displayed.
```

cluster statistics commands

cluster statistics show

Display cluster-wide statistics

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

Description

The `cluster statistics show` command displays the following information. Each item lists the current value and; if applicable, the change (delta) from the previous reported value.

- CPU busy percentage
- Average of CPU busy percentage (advanced privilege level only)
- Total number of NFS and CIFS operations
- Number of NFS operations
- Number of CIFS operations
- Number of cache operations (advanced privilege level only)
- Total amount of network data received (advanced privilege level only)
- Total amount of network data sent (advanced privilege level only)
- Number of packets received (advanced privilege level only)
- Number of packets sent (advanced privilege level only)
- Busy percentage for the data network
- Amount of data received on the data network
- Amount of data sent on the data network
- Busy percentage for the cluster network
- Amount of data received on the cluster network
- Amount of data sent on the cluster network
- Amount of data read from disk
- Amount of data written to disk

At the diagnostic privilege level, the command displays the following information:

- Average of CPU busy percentage
- CPU busy percentage
- Total number of operations
- Number of NFS operations
- Number of CIFS operations
- Number of 7M Fcache operations
- Number of SpinFS operations
- Total amount of network traffic received
- Total amount of network traffic sent
- Percentage of data-network utilization

- Amount of data-network traffic received
- Amount of data-network traffic sent
- Percentage of cluster-network utilization
- Amount of cluster-network traffic received
- Amount of cluster-network traffic sent
- Amount of data read from disk
- Amount of data written to disk
- Number of packets received
- Number of packets sent

Examples

The following example displays cluster statistics:

```
cluster1::> cluster statistics show
      Counter         Value         Delta
------------- ------------- -------------
    CPU Busy:           84%           +27
   Operations:
       Total:     951471448    7210/s:11s
         NFS:   12957951479   13759/s:11s
        CIFS:     342195460     230/s:11s
 Data Network:
        Busy:            0%             -
    Received:        1.98TB  3.18MB/s:11s
        Sent:        6.20TB   903KB/s:11s
Cluster Network:
        Busy:            0%             -
    Received:        6.33TB  1.34MB/s:11s
        Sent:        6.24TB  3.54MB/s:11s
 Storage Disk:
        Read:         207TB  82.7MB/s:11s
       Write:        53.3TB  53.5MB/s:11s
```

cluster time-service commands

cluster time-service ntp key create

Create an NTP symmetric authentication key

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

Description

The `cluster time-service ntp key create` command creates a cryptographic key that can be used to verify that Network Time Protocol (NTP) packets are coming from a configured NTP server.

To use the created key it must be assigned to the required NTP server configuration using the cluster time-service ntp server create or cluster time-service ntp server modify commands.

> The id, key-type and value must all be configured identically on both the ONTAP cluster and the external NTP time server for the cluster to be able to synchronize time to that server.

Parameters

`-id <integer>` - NTP Symmetric Authentication Key ID

Uniquely identifies this key in the cluster. Must be an integer between 1 and 65535.

`-type <sha1>` - NTP Symmetric Authentication Key Type

The cryptographic algorithm that this key is used with. Only SHA1 is currently supported.

`-value <text>` - NTP Symmetric Authentication Key Value

A 40 character hexadecimal digit string that represents a cryptographic key that is shared with the NTP server.

Examples

The following example creates a new SHA-1 NTP symmetric authentication key.

```
cluster1::> cluster time-service ntp key create 1 sha1
2e874852e7d41cda65b23915aa5544838b366c51
```

Related Links

- cluster time-service ntp server create
- cluster time-service ntp server modify

cluster time-service ntp key delete

Delete an NTP symmetric authentication key

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

Description

Delete an NTP symmetric authentication key.

> It is not possible to delete a key that is referenced in an existing NTP server configuration. Remove all references to this key using the cluster time-service ntp server modify or cluster time-service ntp server delete commands before attempting to delete the key using this command.

Parameters

`-id <integer>` - NTP Symmetric Authentication Key ID

Unique identifier of this key in the cluster.

Examples

The following example deletes the NTP key with ID 1.

```
cluster1::> cluster time-service ntp key delete 1
```

Related Links

- cluster time-service ntp server modify
- cluster time-service ntp server delete

cluster time-service ntp key modify

Modify an NTP symmetric authentication key

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

Description

The `cluster time-service ntp key modify` command modifies a Network Time Protocol (NTP) symmetric authentication key.

Parameters

`-id <integer>` - NTP Symmetric Authentication Key ID

Unique identifier of this key in the cluster.

`[-type <sha1>]` - NTP Symmetric Authentication Key Type

The cryptographic algorithm that this key is used with. Only SHA1 is currently supported.

`[-value <text>]` - NTP Symmetric Authentication Key Value

A 40 character hexadecimal digit string that represents a cryptographic key that is shared with the NTP server.

Examples

The following example modifies the NTP key with ID 1 to have a new value.

```
cluster1::> cluster time-service ntp key modify 1 -value
2e874852e7d41cda65b23915aa5544838b366c51
```

cluster time-service ntp key show

Display NTP symmetric authentication keys

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

Description

The `cluster time-service ntp key show` command displays the configured Network Time Protocol (NTP) symmetric authentication keys.

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.

[-id <integer>] - NTP Symmetric Authentication Key ID

If this parameter is specified, the command displays the keys that match the specified key ID.

[-type <sha1>] - NTP Symmetric Authentication Key Type

If this parameter is specified, the command displays the keys that match the specified key type.

[-value <text>] - NTP Symmetric Authentication Key Value

If this parameter is specified, the command displays the keys that match the specified value.

Examples

The following example displays information about the NTP authentication keys in the cluster:

```
cluster1::> cluster time-service ntp key show
ID      Type       Value
------  ---------  ----------------------------------------
2       sha1       5a01120580b5a6ade6ebcd5bad7673fdd6db0113
10      sha1       1f48d2e6f02f17e3f8fa8798cc77af101df29642
2 entries were displayed.
```

cluster time-service ntp security modify

Modify NTP security settings

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

Description

The `cluster time-service ntp security modify` command allows setting of security parameters

related to the Network Time Protocol (NTP) subsystem.

Parameters

`[-is-query-enabled {true|false}]` - Is Querying of NTP Server Enabled? (privilege: advanced)

Setting this parameter to *`true`* allows querying of the NTP subsystem from systems external to the cluster. For example, querying a node using the standard "ntpq" command can be enabled by this command. The default setting is *`false`* to protect against possible security vulnerabilities. If querying of the NTP subsystem is disabled, the cluster time-service ntp status show command can be used to obtain similar information. Although querying of the NTP subsystem from external hosts can be disabled with this command, executing a local query to the localhost address is always enabled.

Examples

The following example enables the querying of the NTP subsystem from clients external to the cluster:

```
cluster-1::> cluster time-service ntp security modify -is-query-enabled
true
```

Related Links

- cluster time-service ntp status show

cluster time-service ntp security show

Display NTP security settings

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

Description

The `cluster time-service ntp security show` command displays the configuration of security features related to the Network Time Protocol (NTP) subsystem.

Examples

The following example displays the NTP security configuration of the cluster:

```
cluster1::> cluster time-service ntp security show
External queries enabled?: true
```

cluster time-service ntp server create

Add a NTP Server

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

Description

The `cluster time-service ntp server create` command associates the cluster with an external network time server for time correction and adjustment by using the Network Time Protocol (NTP).

The command resolves the time server host name to an IP address and performs several validation checks. If an error is detected during validation, it is reported.

The validation checks performed by this command include the following:

- The NTP replies to an NTP query with the specified protocol version.
- The NTP reply indicates that the external time server is synchronized to another time server.
- The distance and dispersion of the NTP reply from the "root" or source clock are within the required limits.

Parameters

`-server <text>` - NTP Server Host Name, IPv4 or IPv6 Address

This parameter specifies the host name or IP address of the external NTP server that is to be associated with the cluster for time correction and adjustment.

`[-version {3|4|auto}]` - NTP Version for Server (default: auto)

Use this parameter to optionally specify the NTP protocol version that should be used for communicating with the external NTP server. If the external NTP server does not support the specified protocol version, time exchange cannot take place.

The supported values for this parameter include the following:

- `3` - Use NTP protocol version 3, which is based on Internet Standard request for comments (RFC) #1305.
- `4` - Use NTP protocol version 4, which is based on Internet Standard RFC #5905.
- `auto` - Have ONTAP select the NTP protocol version.

The default setting is `auto` .

`[-is-preferred {true|false}]` - Is Preferred NTP Server (default: false) (privilege: advanced)

Use this parameter to optionally specify whether the external NTP server is the primary time source for correcting and adjusting the cluster time. The responses from this source are used unless its time is outside the accepted selection range.

The default setting is `false` .

You use this parameter when a high quality radio (or GPS) based time server is being used with a set of non-radio based backup time servers.

`[-key-id <integer>]` - NTP Symmetric Authentication Key ID

Use this parameter to optionally enable NTP symmetric authentication key for communication with the specified time server. The ID must refer to a key created by the cluster time-service ntp key create command and must be a key with the same ID and value as one configured on the specified time server.

Examples

The following example associates the cluster with an NTP server named ntp1.example.com.

```
cluster1::> cluster time-service ntp server create -server
ntp1.example.com
```

Related Links

- cluster time-service ntp key create

cluster time-service ntp server delete

Delete a NTP Server

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

Description

The `cluster time-service ntp server delete` command removes the association between the cluster and an external network time server that uses the Network Time Protocol (NTP).

Parameters

`-server <text>` - NTP Server Host Name, IPv4 or IPv6 Address

This specifies the host name or IP address of an existing external NTP server that the cluster will disassociate from.

Examples

The following example disassociates an NTP server named ntp2.example.com from the cluster:

```
cluster1::> cluster time-service ntp server delete -server
ntp2.example.com
```

cluster time-service ntp server modify

Modify NTP Server Options

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

Description

The `cluster time-service ntp server modify` command modifies the configuration of an existing external network time server that uses the Network Time Protocol (NTP) for time correction and adjustment.

Parameters

`-server <text>` - NTP Server Host Name, IPv4 or IPv6 Address

This parameter specifies the host name or IP address of an existing external NTP server that is to be modified.

`[-version {3|4|auto}]` - NTP Version for Server (default: auto)

Use this parameter to optionally specify the NTP protocol version that should be used for communicating with the external NTP server. If the external NTP server does not support the specified protocol version, time exchange cannot take place.

The supported values for this parameter include the following:

- `3` - Use NTP protocol version 3, which is based on Internet Standard request for comments (RFC) #1305.
- `4` - Use NTP protocol version 4, which is based on Internet Standard RFC #5905.
- `auto` - Have ONTAP select the NTP protocol version.

The default setting is auto.

`[-is-preferred {true|false}]` - Is Preferred NTP Server (default: false) (privilege: advanced)

Use this parameter to optionally specify whether the external NTP server is the primary time source for correcting and adjusting the cluster time. The responses from this source are used unless its time is outside the accepted selection range

The default setting is `false` .

You use this parameter when a high quality radio (or GPS) based time server is being used with a set of non-radio based backup time servers.

This parameter is available only at the advanced privilege level and higher.

`[-is-authentication-enabled {true|false}]` - Is NTP Symmetric Key Authentication Enabled

Use this parameter to optionally disable NTP symmetric key authentication for communication with the specified time server. Using this parameter and selecting *`false`* disables the NTP symmetric key authentication and clears the `key-id` parameter for the specified server. This parameter is not required to enable NTP symmetric key authentication, but if specified as *`true`* the NTP symmetric authentication key must also be specified using the `key-id` parameter.

`[-key-id <integer>]` - NTP Symmetric Authentication Key ID

Use this parameter to optionally enable NTP symmetric authentication key for communication with the specified time server. The ID must refer to a key created by the cluster time-service ntp key create command and must be a key with the same ID and value as one configured on the specified time server.

Examples

The following example modifies the NTP version of an NTP server named ntp1.example.com. The NTP version is changed to 4.

```
cluster1::> cluster time-service ntp server modify -server
ntp1.example.com -version 4
```

Related Links

- cluster time-service ntp key create

cluster time-service ntp server reset

Reset NTP server list to a default selection

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

Description

The `cluster time-service ntp server reset` command replaces the current configuration with one of the selected configurations.

If none or more than one time service configuration is selected, the command will fail.

Parameters

`[-use-public {true|false}]` - Reset Server List to Public Identified Defaults (default: false) (privilege: advanced)

When set to true, this specifies that the public NTP server list used by ONTAP should replace the current configuration.

The default setting is `false` .

Examples

The following example replaces the current time service configuration with the default public NTP server list that is used by ONTAP.

```
cluster1::> cluster time-service ntp server reset -use-public true
```

cluster time-service ntp server show

Display a list of NTP Servers

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

Description

The `cluster time-service ntp server show` command displays the association between the cluster and external network time servers that use the Network Time Protocol (NTP).

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command only displays the fields that you specify. For example: `-fields server, version` .

| [-instance] }

If this parameter is specified, the command displays all the available field information.

[-server <text>] - NTP Server Host Name, IPv4 or IPv6 Address

If this parameter is specified, the command displays the external NTP servers that match the specified server name or IP address.

[-version {3|4|auto}] - NTP Version for Server (default: auto)

If this parameter is specified, the command displays the external NTP servers that use the specified NTP version.

[-is-preferred {true|false}] - Is Preferred NTP Server (default: false) (privilege: advanced)

If this parameter is specified, the command displays the external NTP server or servers that match the specified preferred server status.

[-is-public {true|false}] - Is Public NTP Server Default (privilege: advanced)

If this parameter is specified, the command displays the information for the external NTP servers that are either on the NTP server list defined by ONTAP (`true` ') or not on the list (false`').

[-is-authentication-enabled {true|false}] - Is NTP Symmetric Key Authentication Enabled

If this parameter is specified, the command displays the external NTP server or servers that require NTP symmetric key authentication for communication.

[-key-id <integer>] - NTP Symmetric Authentication Key ID

If this parameter is specified, the command displays the external NTP server or servers that match the specified symmetric authentication key ID.

Examples

The following example displays information about all external NTP time servers that are associated with the cluster:

```
cluster1::> cluster time-service ntp server show
                                          Is
                                          Authentication
Server                          Version   Enabled        Key ID
------------------------------- -------   -------------- ------
ntp1.example.com                auto      false          -
ntp2.example.com                auto      true           10
```

cluster time-service ntp status show

Display status of the node's NTP client

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

Description

The `cluster time-service ntp status show` command displays the status of the associations between the cluster and external network time servers that use the Network Time Protocol (NTP).

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 this parameter is specified, the command displays detailed information about all entries.

`[-node {<nodename>|local}]` - Node (privilege: advanced)

If this parameter is specified, the command displays information related to associations on the specified node.

`[-server <text>]` - NTP Server Host Name, IPv4 or IPv6 Address (privilege: advanced)

If this parameter is specified, the command displays information about the associations related to the specified NTP server. The server should be specified as it is configured in the cluster time-service ntp server show command.

`[-server-address <IP Address>]` - Server IP Address (privilege: advanced)

If this parameter is specified, the command displays information about the associations related to the NTP server with the specified IP address.

`[-is-peer-reachable {true|false}]` - Is Peer Reachable and Responding to Polls? (privilege: advanced)

If this parameter is specified as *`true`* , the command displays information about associations with the NTP servers that have been successfully polled.

`[-is-peer-selected {true|false}]` - Is Peer Selected as Clock Source? (privilege: advanced)

If this parameter is specified as *`true`* , the command displays information about associations with the NTP servers that have been selected as the current clock source.

`[-selection-state <State of NTP Peer Selection>]` - State of Server Selection (privilege: advanced)

If this parameter is specified, the command displays information about associations with the specified selection state.

`[-selection-state-text <text>]` - Description of Server Selection State (privilege: advanced)

If this parameter is specified, the command displays information about associations with the specified selection state description.

`[-poll-interval <integer>]` - Poll Interval (secs) (privilege: advanced)

If this parameter is specified, the command displays information about associations that have the specified polling interval.

[-time-last-poll <integer>] - Time from Last Poll (secs) (privilege: advanced)

If this parameter is specified, the command displays information about associations that are polled at the specified time.

[-offset <double hundredths>] - Offset from Server Time (ms) (privilege: advanced)

If this parameter is specified, the command displays information about associations with the specified offset from the NTP server.

[-delay <double hundredths>] - Delay Time to Server (ms) (privilege: advanced)

If this parameter is specified, the command displays information about associations with the specified travelling time to the NTP server.

[-jitter <double hundredths>] - Maximum Offset Error (ms) (privilege: advanced)

If this parameter is specified, the command displays information about associations with the specified offset error from the NTP server.

[-reachability <Hex String>] - Reachability of Server (privilege: advanced)

If this parameter is specified, the command displays information about associations with the specified reachability to the NTP server. Reachability is specified as a hexbyte that has a bit to represent the success of each of the last eight polls of the specified server. A set bit represents a successful poll. The least significant bit represents the most recent poll, the next most significant bit the poll before that, and so on.

[-server-stratum <integer>] - Stratum of Server Clock (privilege: advanced)

If this parameter is specified, the command displays information about associations with NTP servers with the specified clock stratum.

[-server-reference <text>] - Reference Clock at Server (privilege: advanced)

If this parameter is specified, the command displays information about associations with NTP servers using the specified clock as reference.

[-reported-errors <NTP Peer and Packet Errors>,…] - Reported Packet and Peer Errors (privilege: advanced)

If this parameter is specified, the command displays information about associations with the specified errors.

Examples

The following example displays the status of the NTP associations of the cluster:

```
cluster-1::*>cluster time-service ntp status show
Node: node-1
Server                   Reachable  Selection State                Offset
(ms)
------------------------ ---------  ------------------------------
------------
ntp1.eng.netapp.com          true  Currently Selected Server
39.122
ntp2.eng.netapp.com          true  Candidate Server
37.786
2 entries were displayed.
```

The following example displays the status of the association with the specified external NTP server:

```
cluster-1::*>cluster time-service ntp status show -instance -server
ntp1.example.com
Node: node-1
NTP Server Host Name, IPv4 or IPv6 Address: ntp1.example.com
                         Server IP Address: 10.56.32.33
Is Peer Reachable and Responding to Polls?: true
         Is Peer Selected as Clock Source?: true
                 State of Server Selection: sys_peer
     Description of Server Selection State: Currently Selected Server
                      Poll Interval (secs): 64
                Time from Last Poll (secs): 1
               Offset from Server Time (ms): 26.736
                 Delay Time to Server (ms): 61.772
                 Maximum Offset Error (ms): 3.064
                    Reachability of Server: 01
                   Stratum of Server Clock: 2
                 Reference Clock at Server: 10.56.68.21
           Reported Packet and Peer Errors: -
```

Related Links

- cluster time-service ntp server show

event commands

event catalog commands

event catalog show

Display event definitions

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

Description

The `event catalog show` command displays information about events in the catalog. By default, this command displays the following information:

- Message name of the event
- Severity of the event
- SNMP trap type of the event

To display detailed information about a specific event, run the command with the `-message-name` parameter, and specify the name of the event. The detailed view adds the following information:

- Full description of the event
- Action to be taken to address the event
- Event's deprecation status

You can specify additional parameters to limit output to the information that matches those parameters. For example, to display information only about events with an event name that begins with *`raid`* , enter the command with the `-message-name` *`raid*`* parameter. The parameter value can either be a specific text string or a wildcard pattern.

Alternatively, an event filter can also be specified to limit the output events.

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.

[-message-name <Message Name>] - Message Name

Selects the events that match this parameter value.

[-filter-name <text>] - Filter Name

Selects the events that match this parameter value. The parameter value indicates an existing filter name that, when applied permits the inclusion of the listed events.

[-severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}] - Severity

Selects the events that match this parameter value.

[-description <text>] - Description

Selects the events that match this parameter value.

[-action <text>] - Corrective Action

Selects the events that match this parameter value.

[-snmp-trap-type {Standard|Built-in|Severity-based}] - SNMP Trap Type

Selects the events that match this parameter value. The parameter value describes the type of SNMP trap associated with the event. The value can be one of the following: *Standard* trap type events are those defined in the RFCs. *Built-in* trap types are those that are NetApp Enterprise traps specific to events. The remaining events are considered to have *Severity-based* SNMP trap types.

[-deprecated {true|false}] - Is Deprecated

Selects the events that match this parameter value. The parameter value indicates whether the event is deprecated or not.

> Deprecated events may be removed in a future release of ONTAP.

Examples

The following example displays the event catalog:

```
cluster1::> event filter show -filter-name filter1
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
filter1
            1    include  zapi.*           *             *         *=*
            2    exclude  *                *             *         *=*
2 entries were displayed.

cluster1::> event catalog show -filter-name filter1
Message                           Severity         SNMP Trap Type
--------------------------------- ---------------- -----------------
zapi.killed                       NOTICE           Severity-based
zapi.method.notfound              NOTICE           Severity-based
zapi.sf.up.ready                  INFORMATIONAL    Severity-based
zapi.snapshot.success             NOTICE           Severity-based
zapi.streamout.noMethod           NOTICE           Severity-based
5 entries were displayed.

cluster1::> event catalog show -message-name zsm.* -filter-name filter1
```

```
There are no entries matching your query.

cluster1::> event catalog show -message-name zapi.* -filter-name filter1
Message                           Severity         SNMP Trap Type
--------------------------------- ---------------- -----------------
zapi.method.notfound              NOTICE           Severity-based
zapi.sf.up.ready                  INFORMATIONAL    Severity-based
zapi.snapshot.success             NOTICE           Severity-based
zapi.streamout.noMethod           NOTICE           Severity-based
4 entries were displayed.

cluster1::> event catalog show -message-name CR.*
Message                           Severity         SNMP Trap Type
--------------------------------- ---------------- -----------------
CR.Corrupt.Redir.Deleted          INFORMATIONAL    Severity-based
CR.Dangling.Redir.Deleted         INFORMATIONAL    Severity-based
CR.Data.File.Inaccessible         NOTICE           Severity-based
CR.Del.Corrupt.Redir.Failed       NOTICE           Severity-based
CR.Del.CrptStreamData.Fail        NOTICE           Severity-based
CR.Del.CrptStreamRedir.Fail       NOTICE           Severity-based
CR.Del.DangStreamData.Fail        NOTICE           Severity-based
CR.Del.DangStreamRedir.Fail       NOTICE           Severity-based
CR.Del.Dangling.Redir.Failed      NOTICE           Severity-based
CR.Fix.Corrupt.Redir.Failed       NOTICE           Severity-based
CR.Fix.Crpt.Data.Dir.Failed       INFORMATIONAL    Severity-based
CR.Fix.Crpt.Data.File.Failed      NOTICE           Severity-based
CR.Fix.CrptStreamRedir.Fail       NOTICE           Severity-based
CR.Fix.Dang.Data.File.Failed      NOTICE           Severity-based
CR.Fix.Nlinks.Failed              NOTICE           Severity-based
CR.Fix.TempFiles.Failed           INFORMATIONAL    Severity-based
CR.Max.Session.Exceed             INFORMATIONAL    Severity-based
CR.RDB.Counters.Not.Updated       INFORMATIONAL    Severity-based
CR.RDB.State.Not.Updated          NOTICE           Severity-based
CR.Redir.File.Inaccessible        NOTICE           Severity-based
CR.Snapshot.Not.Deleted           NOTICE           Severity-based

Message                           Severity         SNMP Trap Type
--------------------------------- ---------------- -----------------
CR.Sync.ACL.Fail                  NOTICE           Severity-based
22 entries were displayed.

cluster1::> event catalog show -instance
...
...
     Message Name: Nblade.cifsEncSessAccessDenied
         Severity: ERROR
```

```
      Description: This message occurs when a client not capable of SMB
encryption tries to establish a CIFS session that requires SMB encryption.
Corrective Action: Either ensure that the client is capable of SMB
encryption or disable SMB encryption on the Vserver.
   SNMP Trap Type: Severity-based
    Is Deprecated: false
Message Name: Nblade.cifsEncShrAccessDenied
         Severity: ERROR
      Description: This message occurs when a client not capable of SMB
encryption tries to connect to a CIFS share that requires SMB encryption.
Corrective Action: Either ensure that the client is capable of SMB
encryption or disable SMB encryption on the CIFS share.
   SNMP Trap Type: Severity-based
    Is Deprecated: false
...
```

...

event config commands

event config force-sync

Synchronize a node's EMS configuration with the cluster wide EMS configuration

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

Description

The `event config force-sync` command forces a node's EMS configuration to be synchronized with the cluster wide EMS configuration. The configuration is automatically synchronized among all nodes in the cluster, but in rare cases a node may not be updated. This command simplifies the recovery from this issue.

The following example shows where this command is useful: An email destination is configured for all CRITICAL level event occurrences. When the event is generated, all nodes generate an email except one. This command forces that node to refresh a stale configuration.

Parameters

`[-node {<nodename>|local}]` - Node (privilege: advanced)

The node parameter specifies which controller will be synchronized.

event config modify

Modify log configuration parameters

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

Description

Use the `event config modify` command to configure event notification and logging for the cluster.

Parameters

`[-mail-from <mail address>]` - Mail From

Use this parameter to configure the email address from which email notifications will be sent. You can configure the cluster to send email notifications when specific events occur. Use the `event route add-destinations` and `event destination create` commands to configure email destinations for events.

`[-mail-server <text>]` - Mail Server (SMTP)

Use this parameter to configure the name or IP address of the SMTP server used by the cluster when sending email notification of events. If authentication is required to the mail-server, specify the user name for the mail-server using the *`-mail-server-user`* parameter. Use the event config set-mail-server-password command to set the password used for this user name.

You can optionally specify a port value for the mail server. The port specification for a mail host consists of a colon (":") and a decimal value between 1 and 65335, and follows the mailhost name (for example, *`mymailhost.example.com:5678`*).

`[-mail-server-user <text>]` - Mail Server User Name

If authentication is required to the mail-server, use this parameter to specify the user name for the mail-server specified by the *`-mail-server`* parameter. Use the event config set-mail-server-password command to set the password used for this user name.

`[-suppression {on|off}]` - Event Throttling/Suppression (privilege: advanced)

Use this parameter to configure whether event suppression algorithms are enabled ("on") or disabled ("off"). The event processing system implements several algorithms to throttle duplicate events.

`[-console {on|off}]` - Console Logging (privilege: advanced)

Use this parameter to configure whether events are displayed on the console port ("on") or not displayed("off").

`[-proxy-url <text>]` - HTTP/HTTPS Proxy URL

If your organization uses a proxy, use this parameter to specify an HTTP or HTTPS proxy for REST API type EMS notification destinations. The URL must start with an http:// or https:// prefix. If using an HTTPS proxy, you also need to install the correct Root CA certificates in ONTAP. To specify a URL that contains a question mark, press ESC followed by the "?". Setting this field to an empty string or '-' will clear all proxy settings including the URL, user and password.

`[-proxy-user <text>]` - User Name for HTTP/HTTPS Proxy

If authentication is required, use this parameter to specify the user name for the HTTP or HTTPS proxy server specified by the *`-proxy-url`* parameter. Use the event config set-proxy-password command to set the password used for this user name.

`[-is-pubsub-enabled {true|false}]` - Is Publish/Subscribe Messaging Enabled?

Use this parameter to configure whether or not events are published to the Publish/Subscribe messaging broker.

Examples

The following command sets the "Mail From" address for event notifications to "admin@example.com" and the "Mail Server" to "mail.example.com":

```
cluster1::> event config modify -mailfrom admin@example.com -mailserver
mail.example.com
```

The following command configures a proxy that requires authentication:

```
cluster1::> event config modify -proxy-url http://proxy.example.com:8080
-proxy-user-name admin
cluster1::> event config set-proxy-password

Enter the password:
Confirm the password:
```

The following example turns on event suppression and console logging:

```
cluster1::> event config modify -suppression on -console on
```

Related Links

- event config set-mail-server-password
- event config set-proxy-password

event config set-mail-server-password

Modify password for mail-server

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

Description

Use the `event config set-mail-server-password` command to set the password for authenticated access to a mail-server being used for EMS notifications. This password is used with the user name you specify using the event config modify -mail-server-user command to send EMS messages to email destinations through the mail-server you specify using the event config modify -mail-server command. If you enter the command without parameters, the command prompts you for a password and for a confirmation of that password. Enter the same password at both prompts. The password is not displayed. If you want to clear the mail-server password, use the event config modify -mail-server-user command and set the user-name to an empty string or '-'.

Parameters

Examples

The following example shows successful execution of this command:

```
cluster1::> event config set-mail-server-password

Enter the password:
Confirm the password:
```

Related Links

- event config modify

event config set-proxy-password

Modify password for proxy server

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

Description

Use the `event config set-proxy-password` command to set the password for authenticated access to an HTTP or HTTPS proxy being used for EMS notifications. This password is used with the user name you specify using the event config modify -proxy-user command to send EMS messages to REST API destinations through the proxy you specify using the event config modify -proxy-url command. IF you enter the command without parameters, the command prompts you for a password and for a confirmation of that password. Enter the same password at both prompts. The password is not displayed. If you want to clear the proxy password, use the event config modify -proxy-url command and set the URL to an empty string or '-'.

Parameters

Examples

The following example shows successful execution of this command:

```
cluster1::> event config set-proxy-password

Enter the password:
Confirm the password:
```

Related Links

- event config modify

event config show

Display log configuration parameters

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

Description

The `event config show` command displays information about the configuration of event notification and event logging for the cluster.

"Mail From" is the email address that the event notification system uses as the "From" address for email notifications.

"Mail Server" is the name or IP address of the SMTP server that the event notification system uses to send email notification of events.

"Mail Server User Name" is the user name for the mail-server if authentication is required.

"Proxy URL" is the HTTP or HTTPS proxy server URL used by rest-api type EMS notification destinations if your organization uses a proxy.

"Proxy User Name" is the user name for the HTTP or HTTPS proxy server if authentication is required.

"Is Publish/Subscribe Messaging Enabled?" indicates whether or not events are published to the Publish/Subscribe messaging broker.

"Suppression" indicates whether event suppression algorithms are enabled ("on") or disabled ("off"). The event processing system implements several algorithms to throttle events.

> The suppression parameter can disable both autosuppression and duplicate suppression, but not timer suppression.

"Console" indicates whether events are displayed on the console port ("on") or not displayed("off").

Examples

The following example displays the configuration of event notification for the cluster:

```
cluster1::> event config show
                              Mail From:  admin@example.com
                            Mail Server:  mail.example.com
                              Proxy URL:  -
                        Proxy User Name:  -
    Publish/Subscribe Messaging Enabled:  true
```

The following example displays the configuration of event notification with HTTP or HTTPS proxy:

```
cluster1::> event config show
                              Mail From:  admin@example.com
                            Mail Server:  mail.example.com
                              Proxy URL:  http://proxy.example.com:3128
                        Proxy User Name:  admin
    Publish/Subscribe Messaging Enabled:  true
```

event filter commands

event filter copy

Copy an event filter

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

Description

The `event filter copy` command copies an existing filter to a new filter. The new filter will be created with rules from the source filter. For more information, see the event filter create command.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter to copy.

`-new-filter-name <text>` - New Event Filter Name

Use this mandatory parameter to specify the name of the new event filter to create and copy the rules.

Examples

The following example copies an existing event filter named emer-wafl-events to a new filter named filter1:

```
cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
emer-wafl-events
            1    include  wafl.*           EMERGENCY     *         *=*
            2    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
```

```
                                                                *          *=*
            2    exclude  *                *             *          *=*
10 entries were displayed.

cluster1::> event filter copy -filter-name emer-wafl-events -new-filter
-name filter1

cluster1::> event filter show
Filter      Rule Rule                                      SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                                *          *=*
            2    include  *                *             Standard, Built-
in
                                                                           *=*
            3    exclude  *                *             *          *=*
emer-wafl-events
            1    include  wafl.*           EMERGENCY     *          *=*
            2    exclude  *                *             *          *=*
filter1
            1    include  wafl.*           EMERGENCY     *          *=*
            2    exclude  *                *             *          *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                                *          *=*
            2    include  callhome.*       ERROR         *          *=*
            3    exclude  *                *             *          *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                                *          *=*
            2    exclude  *                *             *          *=*
12 entries were displayed.
```

Related Links

- event filter create

event filter create

Create a new event filter.

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

Description

The `event filter create` command creates a new event filter.

- An event filter is used to select the events of interest and is made up of one or more rules, each of which contains the following three fields:
- name - event (message) name.
- severity - event severity.
- snmp-trap-type - event SNMP trap type.

```
 These fields are evaluated for a match using a logical "AND" operation:
name AND severity AND SNMP trap type. Within a field, the specified
values are evaluated with an implicit logical "OR" operation. So, if `-
snmp-trap-type```_Standard, Built-in_`` is specified, then the event
must match ``_Standard_`` OR ``_Built-in_`` . The wildcard matches all
values for the field.
* Type - include or exclude. When an event matches an include rule, it
will be included into the filter, whereas it will be excluded from the
filter if it matches an exclude rule.
```

Rules are checked in the order they are listed for a filter, until a match is found. There is an implicit rule at the end that matches every event to be excluded. For more information, see the `event filter rule` command.

There are three system-defined event filters provided for your use:

- default-trap-events - This filter matches all ALERT and EMERGENCY events. It also matches all Standard, Built-in SNMP trap type events.
- important-events - This filter matches all ALERT and EMERGENCY events.
- no-info-debug-events - This filter matches all non-INFO and non-DEBUG messages (EMERGENCY, ALERT, ERROR and NOTICE).

The system-defined event filters cannot be modified or deleted.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter to create. An event filter name is 2 to 64 characters long. Valid characters are the following ASCII characters: A-Z, a-z, 0-9, "*", and "-". The name must start and end with: A-Z, a-z, "*", or 0-9.

`[-access-control-role <text>]` - Access Control Role (privilege: advanced)

Use this parameter to specify the access control role of the event filter. Access control role indicates the user role which created the filter and is used to control access to the filter based on RBAC rules.

> This is an optional field. If not specified, the currently logged in user role is used. If created by the 'admin' user, the field is left unset.

Examples

The following example creates an event filter named filter1:

```
cluster1::> event filter create -filter-name filter1

cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
filter1
            1    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
9 entries were displayed.
```

event filter delete

Delete existing event filters

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

Description

The `event filter delete` command deletes an existing event filter, along with all its rules.

The system-defined event filters cannot be deleted.

For more information, see the event filter create command.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter to delete.

Examples

The following example deletes an event filter named filter1:

```
cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
filter1
            1    include  wafl.*           EMERGENCY     *         *=*
            2    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
10 entries were displayed.

cluster1::> event filter delete -filter-name filter1

cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
```

```
            2    include  *                *             Standard, Built-
in
                                                                    *=*
            3    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
8 entries were displayed.
```

Related Links

- event filter create

event filter prepare-for-revert

Deletes unsupported filter or updates unsupported parameter-criteria (parameter-criteria values other than **=**)

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

Description

The `event filter prepare-for-revert` command can be used to remove event filters or update event filter rules that are not supported when the cluster reverts to the previous release. Event filters with rules having a `parameter-criteria` value other than `*=*` are not supported.

Parameters

{ `-delete-unsupported-filters {true|false}` - Clear Unsupported Filters (privilege: advanced)

Use this parameter to delete the event filters that are not supported in the previous release.

| `-update-unsupported-filter-param-criteria {true|false}` - Update Unsupported Filter Parameter Criteria (privilege: advanced) }

Use this parameter to update the event filter rules that are not supported in the previous release to `*=*` .

Examples

The following shows examples of "event filter prepare-for-revert":

```
cluster1::> event filter show
Filter      Rule Rule                                SNMP Trap
Name        Posn Type     Message Name  Severity     Type      Parameters
```

```
----------- ---- -------- ------------- ----------- --------- ----------
default-trap-events
            1    include  *             EMERGENCY, ALERT
                                                    *         *=*
            2    include  *             *           Standard, Built-in
                                                              *=*
            3    exclude  *             *           *         *=*
important-events
            1    include  *             EMERGENCY, ALERT
                                                    *         *=*
            2    include  callhome.*    ERROR       *         *=*
            3    exclude  *             *           *         *=*
no-info-debug-events
            1    include  *             EMERGENCY, ALERT, ERROR, NOTICE
                                                    *         *=*
            2    exclude  *             *           *         *=*
wafl-filter
            1    include  wafl.*        EMERGENCY   *         vol=xyz
            2    exclude  *             *           *         *=*
10 entries were displayed.

cluster1::*> event filter prepare-for-revert -delete-unsupported-filters
true

cluster1::> event filter show
Filter      Rule Rule                                SNMP Trap
Name        Posn Type     Message Name  Severity    Type      Parameters
----------- ---- -------- ------------- ----------- --------- ----------
default-trap-events
            1    include  *             EMERGENCY, ALERT
                                                    *         *=*
            2    include  *             *           Standard, Built-in
                                                              *=*
            3    exclude  *             *           *         *=*
important-events
            1    include  *             EMERGENCY, ALERT
                                                    *         *=*
            2    include  callhome.*    ERROR       *         *=*
            3    exclude  *             *           *         *=*
no-info-debug-events
            1    include  *             EMERGENCY, ALERT, ERROR, NOTICE
                                                    *         *=*
            2    exclude  *             *           *         *=*
8 entries were displayed.
```

event filter rename

Rename an event filter

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

Description

The `event filter rename` command is used to rename an existing event filter.

There are system-defined event filters provided for your use. The system-defined event filters cannot be modified or deleted.

For more information, see the event filter create command.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter to rename.

`-new-filter-name <text>` - New Event Filter Name

Use this mandatory parameter to specify the new name the event filter should be renamed to.

Examples

The following example renames an existing filter named filter1 as emer-wafl-events:

```
cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
filter1
            1    include  wafl.*           EMERGENCY     *         *=*
            2    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
```

```
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
10 entries were displayed.
cluster1::> event filter rename -filter-name filter1 -new-filter-name
emer-wafl-events

cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
emer-wafl-events
            1    include  wafl.*           EMERGENCY     *         *=*
            2    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
10 entries were displayed.
```

Related Links

- event filter create

event filter show-summary

Display event filter summary

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

Description

The `event filter show-summary` command displays a summary of all the event filters. For more details,

use the event filter show command.

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.

[-filter-name <text>] - Filter Name (privilege: advanced)

Selects the event filters that match this parameter value.

[-rule-count <integer>] - Number of Rules (privilege: advanced)

Selects the event filters that match this parameter value.

[-system-defined {true|false}] - System-Defined Filter (privilege: advanced)

Selects the event filters that match this parameter value. System-defined filters are defined by the system and cannot be modified or deleted.

[-access-control-role <text>] - Access Control Role (privilege: advanced)

Selects the event filters that match this parameter value. The access control role indicates the user role that created the filter and is used to control access to the filter based on RBAC rules. For filters created by 'admin', the access control role is empty (indicated by '-').

Examples

The following example displays the event filter summary:

```
cluster-1::*> event filter show-summary
Filter Name         Rule Count  System-Defined Access Control Role
------------------  ----------  -------------- -------------------
default-trap-events
                    4           true           -
important-events    3           true           -
no-info-debug-events
                    2           true           -
test_filter         1           false          test_role
4 entries were displayed.
```

Related Links

- event filter show

event filter show

Display the list of existing event filters.

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

Description

The `event filter show` command displays all the event filters which are configured. An event filter is used to select the events of interest and is made up of one or more rules, each of which contains the following three fields:

*

- name - event (message) name.
- severity - event severity.
- snmp-trap-type - event SNMP trap type.

```
 These fields are evaluated for a match using a logical "AND" operation:
name AND severity AND SNMP trap type. Within a field, the specified
values are evaluated with an implicit logical "OR" operation. So, if `-
snmp-trap-type```_Standard, Built-in_`` is specified, then the event
must match ``_Standard_`` OR ``_Built-in_`` . The wildcard matches all
values for the field.
* Type - include or exclude. When an event matches an include rule, it
will be included into the filter, whereas it will be excluded from the
filter if it matches an exclude rule.
```

Rules are checked in the order they are listed for a filter, until a match is found. There is an implicit rule at the end that matches every event to be excluded. For more information, see `event filter rule` command.

There are three system-defined event filters provided for your use:

- default-trap-events - This filter matches all ALERT and EMERGENCY events. It also matches all Standard, Built-in SNMP trap type events.
- important-events - This filter matches all ALERT and EMERGENCY events.
- no-info-debug-events - This filter matches all non-INFO and non-DEBUG messages (EMERGENCY, ALERT, ERROR and NOTICE).

The system-defined event filters cannot be modified or deleted.

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.

[-filter-name <text>] - Filter Name

Selects the event filters that match this parameter value.

[-position <integer>] - Rule Position

Selects the event filters that match this parameter value.

[-type {include|exclude}] - Rule Type

Selects the event filters that match this parameter value. The rule types are as follows:

- include - Events matching this rule are included in the specified filter.
- exclude - Events matching this rule are excluded in the specified filter.

[-message-name <text>] - Message Name

Selects the event filters that match this parameter value.

[-severity <text>,…] - Severity

Selects the events that match this parameter value. Severity levels:

- EMERGENCY - Disruption.
- ALERT - Single point of failure.
- ERROR - Degradation.
- NOTICE - Information.
- INFORMATIONAL - Information.
- DEBUG - Debug information.
- * - Includes all severities.

[-snmp-trap-type <text>,…] - SNMP Trap Type

Selects the event filters that match this parameter value. The SNMP trap types are as follows:

- Standard - Traps defined in RFCs.
- Built-in - Enterprise traps specific to events.
- Severity-based - Traps specific to events that do not belong to the above two types.
- * - Includes all SNMP trap types.

[-parameter-criteria [key>=<value],…] - Parameter Criteria

Selects the event filters that match this parameter-criteria value.

[-system-defined {true|false}] - System-Defined Filter

Selects the event filters that match this parameter value.

[-access-control-role <text>] - Access Control Role (privilege: advanced)

Selects the event filters that match this parameter value.

Examples

The following example displays the event filters:

```
cluster1::> event filter show
Filter      Rule Rule                                      SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    include  *                *             Standard, Built-
in
                                                                    *=*
            4    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                          *         *=*
            2    exclude  *                *             *         *=*
9 entries were displayed.
```

The following example displays the event filters queried on the SNMP trap type value "Standard":

```
cluster1::> event filter show -snmp-trap-type Standard
Filter      Rule Rule                                      SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            3    include  *                *             Standard, Built-
in
                                                                    *=*
```

The following example displays the event filters with one or more rules that have no condition on the SNMP trap type. Note that the wildcard character has to be specified in double-quotes. Without double-quotes, output would be the same as not querying on the field.

```
cluster1::> event filter show -snmp-trap-type "*"
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  callhome.*       ERROR         *         *=*
            4    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                          *         *=*
            2    exclude  *                *             *         *=*
8 entries were displayed.
```

event filter test

Test an event filter

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

Description

The `event filter test` command is used to test an event filter. When specified with a message name, the command displays whether the message name is included or excluded from the filter. When specified without a message name, the command displays the number of events from the catalog that match the filter. For more information, see the event filter create command.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter to test.

`[-message-name <Message Name>]` - Message Name

Use this optional parameter to specify the message name of the event to test against the filter.

Examples

The following example tests an event filter named err-wafl-no-scan-but-clone:

```
cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
err-wafl-no-scan-but-clone
            1    include  wafl.scan.clone.*
                                          *             *         *=*
            2    exclude  wafl.scan.*      *             *         *=*
            3    include  wafl.*           EMERGENCY, ALERT, ERROR
                                                         *         *=*
            4    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*

Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
no-info-debug-events
            2    exclude  *                *             *         *=*
12 entries were displayed.

cluster1::> event filter test -filter-name err-wafl-no-scan-but-clone
271 events will be included in the given filter.

cluster1::> event filter test -filter-name err-wafl-no-scan-but-clone
-message-name wafl.scan.clone.split.cantLock
The message-name "wafl.scan.clone.split.cantLock" is included in the given
filter.
```

```
cluster1::> event filter test -filter-name err-wafl-no-scan-but-clone
-message-name wafl.scan.layout.cantWrite
The message-name "wafl.scan.layout.cantWrite" is excluded from the given
filter.
```

Related Links

- event filter create

event filter update-access-control-role

Update access-control-role of an event filter

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

Description

The `event filter update-access-control-role` command is used to update the 'access-control-role' field of an existing event filter.

Parameters

`-filter-name <text>` - Filter Name (privilege: advanced)

Specify the event filter name with this mandatory parameter.

`-new-access-control-role <text>` - New Access Control Role (privilege: advanced)

Specify the new access control role with this mandatory parameter.

Examples

This example shows how to update the access control role of an event filter named filter1:

```
cluster1::*> event filter show-summary
Filter Name         Rule Count  System-Defined Access Control Role
------------------  ----------  -------------- -------------------
default-trap-events
                    4           true           -
filter1             2           false          -
important-events    3           true           -
no-info-debug-events
                    2           true           -
4 entries were displayed.

cluster1::*> event filter update-access-control-role -filter-name filter1
-new-access-control-role new_role

cluster1::*> event filter show-summary
Filter Name         Rule Count  System-Defined Access Control Role
------------------  ----------  -------------- -------------------
default-trap-events
                    4           true           -
filter1             2           false          new_role
important-events    3           true           -
no-info-debug-events
                    2           true           -
4 entries were displayed.
```

event filter rule add

Add a rule for an event filter

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

Description

The `event filter rule add` command adds a new rule to an existing event filter. See event filter create for more information on event filters and how to create a new event filter.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter to add the rule. Rules cannot be added to system-defined event filters.

`[-position <integer>]` - Rule Position

Use this optional parameter to specify the position of the rule in the event filter. It should be in the range (1..n-1), where 'n' is the position of the last rule, which is an implicit rule. Rules are checked in the order they are listed for a filter, until a match is found.

-type {include|exclude} - Rule Type

Use this mandatory parameter to specify the type of the rule which determines whether to include or exclude the events that match this rule.

[-message-name <text>] - Message Name

Use this parameter to specify the message name of the event to include or exclude from the filter.

[-severity <text>,…] - Severity

Use this parameter to specify the list of severity values to match against the events. Enter multiple severities separated by a comma. To enter all severities, the wild card (*) can be used. The wild card cannot be specified with other severities. The default value is *.

[-snmp-trap-type <text>,…] - SNMP Trap Type

Use this parameter to specify the list of the SNMP trap type values to match against the events. Enter multiple SNMP trap types separated by comma. To enter all SNMP trap types, the wild card (*) can be used. The wild card cannot be specified with other SNMP trap types. The default value is *.

[-parameter-criteria [key>=<value],…] - Parameter Criteria

Use this parameter to match against event parameters. Each parameter consists of a name and a value. When multiple parameter criteria are provided in a rule, they all need to match for the rule to be considered matched. A pattern can include one or more wildcard '*' characters.

Examples

The following example adds a rule to an existing event filter "emer-and-wafl": All events with severity EMERGENCY and message name starting with "wafl.**" are included in the filter. Not specifying the SNMP trap type implies a default value of ""**.

```
cluster1::> event filter rule add -filter-name emer-and-wafl -type include
-message-name wafl.*  -severity EMERGENCY
cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
emer-and-wafl
            1    include  wafl.*           EMERGENCY     *         *=*
            2    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
10 entries were displayed.
```

The following example adds a rule to the event filter "emer-and-wafl" at position 1: All events with severity ALERT and message name starting with "wafl.scan.*" are included in the filter.

```
cluster1::> event filter rule add -filter-name emer-and-wafl -type include
-message-name wafl.scan.* -position 1 -severity ALERT

cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
emer-and-wafl
            1    include  wafl.scan.*      ALERT         *         *=*
            2    include  wafl.*           EMERGENCY     *         *=*
            3    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
11 entries were displayed.
```

The following example adds a rule to the event filter "emer-and-wafl" to include all "Standard" SNMP trap type events:

```
cluster1::> event filter rule add -filter-name emer-and-wafl -type include
-snmp-trap-type Standard

cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
emer-and-wafl
            1    include  wafl.scan.*      ALERT         *         *=*
            2    include  wafl.*           EMERGENCY     *         *=*
            3    include  *                *             Standard  *=*
            4    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
12 entries were displayed.
```

The following example adds a rule to the event filter "emer-and-wafl" to include all "wafl" events whose parameters have a parameter named "type" and its value matches "volume":

```
cluster1::> event filter rule add -filter-name emer-and-wafl -type include
-message-name wafl.* -position 1 -parameter-criteria type=volume

cluster1::> event filter show -filter-name emer-and-wafl
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
emer-and-wafl
            1    include  wafl.*           *             *
type=volume
            2    include  wafl.scan.*      ALERT         *         *=*
            3    include  wafl.*           EMERGENCY     *         *=*
            4    include  *                *             Standard  *=*
            5    exclude  *                *             *         *=*
5 entries were displayed.
```

Related Links

- event filter create

event filter rule delete

Delete a rule for an event filter

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

Description

The `event filter rule delete` command deletes a rule from an event filter. The position of all the rules following the deleted rule is updated to maintain a contiguous sequence. Use event filter show command to view the filters and the rules associated with them.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter from which you want to delete the rule. Rules cannot be deleted from system-defined filters.

`-position <integer>` - Rule Position

Use this mandatory parameter to specify the position of the rule to delete from the filter. It should be in the range (1..n-1), where 'n' is the position of the last rule, which is an implicit rule.

Examples

The following example deletes a rule at position 2 from an existing event filter "emer-and-wafl":

```
cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
emer-and-wafl
            1    include  wafl.scan.*      ALERT         *         *=*
            2    include  wafl.*           EMERGENCY     *         *=*
            3    include  *                *             Standard  *=*
            4    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
12 entries were displayed.
cluster1::> event filter rule delete -filter-name emer-and-wafl -position
2

cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  *                *             Standard, Built-
in
                                                                   *=*
            3    exclude  *                *             *         *=*
emer-and-wafl
            1    include  wafl.scan.*      ALERT         *         *=*
```

```
            2    include  *                *             Standard  *=*
            3    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                         *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                         *         *=*
            2    exclude  *                *             *         *=*
11 entries were displayed.
```

Related Links

- event filter show

event filter rule reorder

Modify the index of a rule for an event filter

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

Description

The `event filter rule reorder` command moves a rule to a new position in an existing event filter. Use event filter show command to display all the event filters and the rules associated with them.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter from which you want to change the position of the rule. Rules from system-defined event filters cannot be modified.

`-position <integer>` - Rule Positon

Use this mandatory parameter to specify the position of the rule you want to change. It should be in the range (1..n-1), where 'n' is the position of the last rule, which is an implicit rule.

`-to-position <integer>` - New Rule Position

Use this mandatory parameter to specify the new position to move the rule. It should be in the range (1..n-1), where 'n' is the position of the last rule, which is an implicit rule.

Examples

The following example changes the position of a rule from 1 to 2 from an existing event filter "emer-and-wafl":

```
cluster1::> event filter show
Filter      Rule Rule                                    SNMP Trap
Name        Posn Type     Message Name     Severity      Type
```

```
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  *                *             Standard, Built-
in
                                                                    *=*
            3    exclude  *                *             *         *=*
emer-and-wafl
            1    include  wafl.scan.*      ALERT         *         *=*
            2    include  *                *             Standard  *=*
            3    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                          *         *=*
            2    exclude  *                *             *         *=*
11 entries were displayed.

cluster1::> event filter rule reorder -filter-name emer-and-wafl -position
1 -to-position 2

cluster1::> event filter show
Filter      Rule Rule                                     SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  *                *             Standard, Built-
in
                                                                    *=*
            3    exclude  *                *             *         *=*
emer-and-wafl
            1    include  *                *             Standard  *=*
            2    include  wafl.scan.*      ALERT         *         *=*
            3    exclude  *                *             *         *=*
important-events
```

```
            1    include  *                EMERGENCY, ALERT
                                                  *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                  *         *=*
            2    exclude  *                *             *         *=*
11 entries were displayed.
```

Related Links

- event filter show

event log commands

event log show

Display latest log events

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

Description

The `event log show` command displays the contents of the event log, which lists significant occurrences within the cluster. Use the event catalog show command to display information about events that can occur.

By default, the command displays EMERGENCY, ALERT and ERROR severity level events with the following information, with the most recent events listed first:

- The time at which the event occurred
- The node on which the event occurred
- The severity of the event
- The event's message

To display detailed information about events, use one or more of the optional parameters that affect how the command output is displayed and the amount of detail that is included. For example, to display all detailed event information, use the `-detail` parameter.

To display NOTICE, INFORMATIONAL or DEBUG severity level events, use the `-severity` parameter.

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.

| [-detail]

Displays additional event information such the sequence number of the event.

| [-detailtime]

Displays detailed event information in reverse chronological order.

| [-instance] }

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

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

Displays a list of events for the node you specify. Use this parameter with the `-seqnum` parameter to display detailed information.

[-seqnum <Sequence Number>] - Sequence#

Selects the events that match this parameter value. Use with the `-node` parameter to display detailed information.

[-time <MM/DD/YYYY HH:MM:SS>] - Time

Selects the events that match this parameter value. Use the format: MM/DD/YYYY HH:MM:SS [+- HH:MM]. You can specify a time range by using the ".." operator between two time statements.

```
show -time "08/13/2010 05:55:00".."08/13/2010 06:10:00"
```

Comparative time values are relative to "now". For example, to display only events that occurred within the last minute:

```
show -time >1m
```

+
NOTE: The month and date fields of this parameter are not zero-padded. These fields can be single digits: for example, "7/1/2019 05:55:00".

+

[-severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}] - Severity

Selects the events that match this parameter value. Severity levels are as follows:

- EMERGENCY - Disruption.
- ALERT - Single point of failure.
- ERROR - Degradation.
- NOTICE - Information.
- INFORMATIONAL - Information.
- DEBUG - Debug information.

To display all events, including ones with severity levels of NOTICE, INFORMATIONAL and DEBUG, specify severity as follows:

```
show -severity <=DEBUG
```

[-ems-severity {NODE_FAULT|SVC_FAULT|NODE_ERROR|SVC_ERROR|WARNING|NOTICE|INFO|DEBUG|VAR}] - EMS Severity (privilege: advanced)

Selects the events that match this parameter value. Severity levels:

- NODE_FAULT - Data corruption has been detected or the node is unable to provide client service
- SVC_FAULT - A temporary loss of service, typically a transient software fault, has been detected
- NODE_ERROR - A hardware error that is not immediately fatal has been detected
- SVC_ERROR - A software error that is not immediately fatal has been detected
- WARNING - A high-priority message that does not indicate a fault
- NOTICE - A normal-priority message that does not indicate a fault
- INFO - A low-priority message that does not indicate a fault
- DEBUG - A debugging message
- VAR - A message with variable severity, selected at runtime.

[-source <text>] - Source

Selects the events that match this parameter value (typically a software module).

[-message-name <Message Name>] - Message Name

Selects the events that match this parameter value (string). Message names are descriptive, so filtering output by message name displays messages of a specific type.

[-event <text>] - Event

Selects the events that match this parameter value. The "event" field contains the full text of the event, including any parameters. For example, a wafl.vol.offline event will contain the name of the volume taken offline.

[-kernel-generation-num <integer>] - Kernel Generation Number (privilege: advanced)

Selects the events that match this parameter value. Only events that emanate from the kernel have kernel generation numbers.

[-kernel-sequence-num <integer>] - Kernel Sequence Number (privilege: advanced)

Selects the events that match this parameter value. Only events that emanate from the kernel have kernel sequence numbers.

[-action <text>] - Corrective Action

Selects the events that match this parameter value. The "action" field describes what steps, if any, you must take to remedy the situation.

[-description <text>] - Description

Selects the events that match this parameter value. The "description" field describes why the event was encountered and what it means.

[-filter-name <Filter Name>] - Filter Name

Selects the events that match this parameter value. Only events that were included by existing filters that match this value are displayed.

Examples

The following example displays the event log:

```
cluster1::> event log show
Time                Node             Severity      Event
------------------- ---------------- -------------
------------------------
11/9/2015 13:54:19  node1            NOTICE        vifmgr.portup: A link
up event was received on node node1, port e0a.
11/9/2015 13:54:19  node1            NOTICE        vifmgr.portup: A link
up event was received on node node1, port e0d.
11/9/2015 13:54:19  node1            NOTICE        vifmgr.portup: A link
up event was received on node node1, port e0c.
11/9/2015 13:54:19  node1            NOTICE        vifmgr.portup: A link
up event was received on node node1, port e0b.
...
```

This example demonstrates how to use a range with the `-time` parameter to display all events that occurred during an extended time period. It displays all events that occurred between 1:45pm and 1:50pm on November 9, 2010.

```
cluster1::> event log show -time "11/9/2015 13:45:00".."11/9/2015 13:50:0"
```

The `-time` parameter also accepts values that are relative to "now". The following example displays events that occurred more than one hour ago:

```
cluster1::event log> show -time <1h
Time                Node             Severity      Event
------------------- ---------------- -------------
------------------------
11/9/2015 13:02:03  node1            INFORMATIONAL
monitor.globalStatus.ok: The system's global status is normal.
11/9/2015 13:02:03  node2            INFORMATIONAL
monitor.globalStatus.ok: The system's global status is normal.
...
```

Severity levels sort in the order opposite to what you might expect. The following example displays all events that have a severity level of ERROR or more severe:

```
cluster1::> event log show -severity <ERROR
```

Related Links

- event catalog show

event notification commands

event notification create

Create an event notification

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

Description

The ` event notification create` command is used to create a new notification of a set of events defined by an event filter to one or more notification destinations.

Parameters

`-filter-name <text>` - Filter Name

Use this mandatory parameter to specify the name of the event filter. Events that are included in the event filter are forwarded to the destinations specified in the destinations parameter.

The filter name passed to this command must be an existing filter. For more information, see the event filter create command.

`-destinations <text>,…` - List of Event Notification Destinations

Use this mandatory parameter to specify the list of destinations to which the notification should be forwarded. Enter multiple destinations separated by a comma.

The destination passed to this command must be an existing destination. For more information, see the `event destination create` command.

`[-access-control-role <text>]` - Access Control Role (privilege: advanced)

Use this parameter to specify the access control role of the event notification. Access control role indicates the user role that created the notification and is used to control access to the notification based on RBAC rules.

> This is an optional field. If not specified, the currently logged in user role is used. If created by the 'admin' user, the field is left unset.

Examples

The following example creates an event notification for filter name "filter1" to destinations "email_dest, snmp-traphost and syslog_dest":

```
cluster1::> event notification destination show
                             Hide
Name            Type         Params   Destination
--------------  ----------   ------   ---------------------
email_dest      email        false    test@example.com
snmp-traphost   snmp         true     10.27.12.1 (from "system snmp
traphost")
syslog_dest     syslog       false    10.23.12.1
3 entries were displayed.

cluster1::> event filter show -filter-name filter1
Filter      Rule Rule                                      SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
filter1
            1    exclude  callhome.bad.ram *              *        *=*
            2    include  callhome.*       ALERT, ERROR   *        *=*
            3    exclude  *                *              *        *=*
3 entries were displayed.

cluster1::> event notification create -filter-name filter1 -destinations
email_dest,syslog_dest,snmp-traphost

cluster1::> event notification show
ID     Filter Name       Destinations
-----  ----------------  -----------------
1      filter1           email_dest, syslog_dest, snmp-traphost
```

Related Links

- event filter create

event notification delete

Delete event notifications

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

Description

The `event notification delete` command deletes an existing event notification.

Parameters

`-ID <integer>` - Event Notification ID

Use this parameter to specify the ID of the notification to be deleted.

Examples

The following example shows the deletion of event notification with ID 1:

```
cluster1::> event notification show
ID     Filter Name       Destinations
-----  ----------------  -----------------
1      filter1           email_dest, syslog_dest, snmp-traphost

cluster1::> event notification delete -ID 1

cluster1::> event notification show
This table is currently empty.
```

event notification modify

Modify event notifications

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

Description

The `event notification modify` command is used to modify an existing notification.

Parameters

`-ID <integer>` - Event Notification ID

Use this mandatory parameter to specify the ID of the notification to be modified.

`[-filter-name <text>]` - Event Filter Name

Use this parameter to specify the filter name to be modified.

`[-destinations <text>,…]` - List of Event Notification Destinations

Use this parameter to specify the destinations to be modified. Enter multiple destinations separated by a comma.

Provide the complete set of destinations to be modified. Individual destinations cannot be added or removed.

Examples

The following example shows the modification of the event notification with ID 1:

```
cluster1::> event notification show
ID     Filter Name       Destinations
-----  ----------------  -----------------
1      filter1           email_dest, syslog_dest, snmp-traphost

cluster1::> event notification modify -ID 1 -destinations email_dest,
syslog_dest

cluster1::>  event notification show
ID     Filter Name       Destinations
-----  ----------------  -----------------
1      filter1           email_dest, syslog_dest
```

event notification show

Display event notifications

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

Description

The `event notification show` command is used to display the list of existing event notifications.

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.

[-ID <integer>] - Event Notification ID

Use this parameter to display the detailed information about the notification ID you specify.

[-filter-name <text>] - Event Filter Name

Use this parameter to display event notifications that use the filter-name you specify.

[-destinations <text>,…] - List of Event Notification Destinations

Use this parameter to display event notifications that use the destinations you specify.

[-access-control-role <text>] - Access Control Role (privilege: advanced)

Use this parameter to display event notifications that use the specified access control role.

Examples

The following example displays the event notification:

```
cluster1::> event notification show
ID     Filter Name       Destinations
-----  ----------------  -----------------
1      filter1           email_dest, syslog_dest, snmp-traphost
```

event notification destination create

Create an event notification destination

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

Description

The `event notification destination create` command creates a new event notification destination of either email or syslog type.

The following system-defined notification destination is configured for your use:

- snmp-traphost - This destination reflects the configuration in "system snmp traphost".

Parameters

`-name <text>` - Destination Name

Use this mandatory parameter to specify the name of the notification destination that is to be created. An event notification destination name must be 2 to 64 characters long. Valid characters are the following ASCII characters: A-Z, a-z, 0-9, "_", and "-". The name must start and end with: A-Z, a-z, or 0-9.

{ `-email <mail address>` - Email Destination

Use this parameter to specify the email address event notifications are sent to. For events to properly generate email notifications, the event system must also be configured with an address and mail server from which the mail will be sent. See event config modify command for more information.

| `-syslog <text>` - Syslog Destination

Use this parameter to specify the syslog server host name or IP address syslog messages are sent to.

`[-syslog-port <integer>]` - Syslog Port

Use this parameter to specify the syslog server port value syslog messages are sent to. The default port used depends on the `syslog-transport` value. If the `syslog-transport` is set to *`tcp-encrypted`*, the `syslog-port` has the default value 6514. If the `syslog-transport` is set to *`tcp-unencrypted`*, the `syslog-port` has the default value 601. Otherwise, the default `syslog-port` is set to 514.

`[-syslog-transport {udp-unencrypted|tcp-unencrypted|tcp-encrypted}]` - Syslog Transport

Use this parameter to specify the transport protocol that is used to send the syslog messages.

The `syslog-transport` can be one of the following values:

- *`udp-unencrypted`* - User Datagram Protocol with no security
- *`tcp-unencrypted`* - Transmission Control Protocol with no security

- *tcp-encrypted* - Transmission Control Protocol with Transport Layer Security (TLS)

The default protocol is *udp-unencrypted* . + If *tcp-encrypted* transport is specified, then ONTAP verifies the identity of the destination host by validating its certificate. If the Online Certificate Status Protocol (OCSP) is enabled for Event Management System (EMS), then ONTAP uses that protocol to determine the certificate's revocation status. Use the (privilege: advanced) security config ocsp show -application ems command to determine if the OCSP-based certificate revocation status check is enabled for EMS.

[-syslog-message-format {legacy-netapp|rfc-5424}] - Syslog Message Format

Use this parameter to specify the message format to be used for EMS syslog messages.

The `syslog-message-format` can be one of the following values:

- *legacy-netapp* - Variation of RFC-3164 Syslog format (format: <PRIVAL>TIMESTAMP [HOSTNAME:Event-name:Event-severity]: MSG)
- *rfc-5424* - Syslog format as per RFC-5424 (format: <PRIVAL>VERSION TIMESTAMP HOSTNAME Event-source - Event-name - MSG)

Refer to the respective RFCs for detailed information on the syslog message formats. + The default message format is *legacy-netapp* .

[-syslog-timestamp-format-override {no-override|rfc-3164|iso-8601-utc|iso-8601-local-time}] - Syslog Timestamp Format Override

Use this parameter to override the default timestamp format (based on the `syslog-message-format` parameter) used for EMS syslog messages.

The `syslog-timestamp-format-override` can be one of the following values:

- *no-override* - Timestamp format based on the `syslog-message-format` parameter (*rfc-3164* if message format is *legacy-netapp* , *iso-8601-local-time* if message format is *rfc-5424*)
- *rfc-3164* - Timestamp format as per RFC-3164 (format: Mmm dd hh:mm:ss)
- *iso-8601-utc* - Timestamp format as per ISO-8601 in UTC (format: YYYY-MM-DDThh:mm:ssZ)
- *iso-8601-local-time* - Timestamp format as per ISO-8601 in local time (format: YYYY-MM-DDThh:mm:ss+/-hh:mm)

The default value is *no-override* . When this parameter is modified, its value persists even when `syslog-message-format` is updated. +

[-syslog-hostname-format-override {no-override|fqdn|hostname-only}] - Syslog Hostname Format Override

Use this parameter to override the default hostname format (based on the `syslog-message-format` parameter) used for EMS syslog messages.

The `syslog-hostname-format-override` can be one of the following values:

- *no-override* - Hostname format based on the `syslog-message-format` parameter (*fqdn* if message format is *rfc-5424* , *hostname-only* if message format is *legacy-netapp*)
- *fqdn* - Fully Qualified Domain Name (e.g., myhost.example.com)

- *hostname-only* - Hostname only, without the domain name (e.g., myhost)

The default value is *no-override* . When this parameter is modified, its value persists even when `syslog-message-format` is updated. +

| -rest-api-url <text> - REST API Server URL

Use this parameter to specify the REST API server URL to which event notifications are sent. Enter the full URL, which must start either with an http:// or https:// prefix. To specify a URL that contains a question mark, press ESC followed by the "?". + If a https:// URL is specified, then ONTAP verifies the identity of the destination host by validating its certificate. If the Online Certificate Status Protocol (OCSP) is enabled for Event Management System (EMS), then ONTAP uses that protocol to determine the certificate's revocation status. Use the (privilege: advanced) security config ocsp show -application ems command to determine if the OCSP-based certificate revocation status check is enabled for EMS.

[-certificate-authority <text>] - Client Certificate Issuing CA

Use this parameter to specify the name of the certificate authority (CA) that signed the client certificate that will be sent in case mutual authentication with the REST API server is required. + There can be multiple client certificates installed for the admin vserver in the cluster, and this parameter, along with the *certificate-serial* parameter, uniquely identifies which one. + Use the security certificate show command to see the list of certificates installed in the cluster.

[-certificate-serial <text>] - Client Certificate Serial Number }

Use this parameter to specify the serial number of the client certificate that will be sent in case mutual authentication with the REST API server is required.

[-access-control-role <text>] - Access Control Role (privilege: advanced)

Use this parameter to specify the access control role of the event notification destination. Access control role indicates the user role which created the destination and is used to control access to the destination based on RBAC rules.

> This is an optional field. If not specified, the currently logged in user role is used. If created by the 'admin' user, the field is left unset.

Examples

The following example shows the creation of a new event notification destination of type email called "StorageAdminEmail":

```
cluster1::> event notification destination create -name StorageAdminEmail
-email StorageAdmin@example.com

cluster1::> event notification destination show

Name            Type        Destination
--------------  ----------  ---------------------
StorageAdminEmail
                email       StorageAdmin@example.com
snmp-traphost   snmp        10.30.40.10 (from "system snmp traphost")
2 entries were displayed.
```

The following example shows the creation of a new event notification destination of type rest-api called "RestApi":

```
cluster1::> event notification destination create -name RestApi -rest-api
-url https://rest.example.com/rest
-certificate-authority cluster1-root-ca -certificate-serial 052213E60B7088

cluster1::> event notification destination show -name RestApi -instance
Destination Name: RestApi
             Type of Destination: rest-api
              Destination Values: https://rest.example.com/rest
   Client Certificate Issuing CA: cluster1-root-ca
Client Certificate Serial Number: 052213E60B7088
       Client Certificate Valid?: -
                     Syslog Port: -
                Syslog Transport: -
           Syslog Message Format: -
Syslog Timestamp Format Override: -
 Syslog Hostname Format Override: -
      System-Defined Destination: false
```

Related Links

- event config modify
- security config ocsp show
- security certificate show

event notification destination delete

Delete existing event destinations

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

Description

The `event notification destination delete` command deletes an event notification destination.

The following system-defined notification destination is configured for your use:

- snmp-traphost - This destination reflects the configuration in "system snmp traphost". To remove snmp-traphost addresses, use the `system snmp traphost` command.

Parameters

`-name <text>` - Destination Name

Use this mandatory parameter to specify the name of an event destination to be removed.

Examples

The following shows the examples of deleting event notification destinations:

```
cluster1::> event notification destination show
Name            Type        Destination
    --------------  ----------  ---------------------
    StorageAdminEmail
                    email       StorageAdmin@example.com
    StorageAdminSyslog
                    syslog      example.com
    snmp-traphost   snmp        10.30.40.10 (from "system snmp traphost")
    3 entries were displayed.
cluster1::> event notification destination delete -name StorageAdminEmail

    cluster1::> event notification destination show

    Name            Type        Destination
    --------------  ----------  ---------------------
    StorageAdminSyslog
                    syslog      example.com
    snmp-traphost   snmp        10.30.40.10 (from "system snmp traphost")
    2 entries were displayed.
cluster1::> event notification destination delete -name Storage*
cluster1::> event notification destination show
Name            Type        Destination
    --------------  ----------  ---------------------
    snmp-traphost   snmp        10.30.40.10 (from "system snmp traphost")
    1 entries were displayed.
```

event notification destination modify

Modify an event notification destination

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

Description

The `event notification destination modify` command modifies an event notification destination. More detailed information about parameters can be found in the man page for the event notification destination create command.

The following system-defined notification destination is configured for your use:

- snmp-traphost - This destination reflects the configuration in "system snmp traphost". To modify traphost addresses, use the `system snmp traphost` command.

Parameters

`-name <text>` - Destination Name

Use this mandatory parameter to specify the name of an event notification destination to be modified. The name of the destination must already exist.

{ `[-email <mail address>]` - Email Destination

Use this parameter to specify a new value of email address to replace the current address in the event notification destination. The parameter is specified only when the event notification destination type is already "email". It is not allowed to specify the parameter for a destination that already has another type of destination address.

| `[-syslog <text>]` - Syslog Destination

Use this parameter to specify a new syslog server host name or IP address to replace the current address of the event notification destination. The parameter is specified only when the event notification destination type is already "syslog". It is not allowed to specify the parameter for a destination that already has another type of destination address.

`[-syslog-port <integer>]` - Syslog Port

Use this parameter to specify a new syslog server port value to replace the current port value of the event notification destination. The parameter is specified only when the event notification destination type is already "syslog". It is not allowed to specify the parameter for a destination that already has another type of destination address.

`[-syslog-transport {udp-unencrypted|tcp-unencrypted|tcp-encrypted}]` - Syslog Transport

Use this parameter to specify a new syslog transport to replace the current transport of the event notification destination. The parameter is specified only when the event notification destination type is already "syslog". It is not allowed to specify the parameter for a destination that already has another type of destination address.

`[-syslog-message-format {legacy-netapp|rfc-5424}]` - Syslog Message Format

Use this parameter to specify a new syslog message format to replace the current message format of the event notification destination.

`[-syslog-timestamp-format-override {no-override|rfc-3164|iso-8601-utc|iso-8601-local-time}]` - Syslog Timestamp Format Override

Use this parameter to override the default syslog timestamp format (based on the `syslog-message-format` parameter) of the event notification destination.

`[-syslog-hostname-format-override {no-override|fqdn|hostname-only}]` - Syslog Hostname Format Override

Use this parameter to override the default syslog hostname format (based on the `syslog-message-format` parameter) of the event notification destination.

| `[-rest-api-url <text>]` - REST API Server URL

Use this parameter to specify a new REST API server URL to replace the current address of the event notification destination. Enter the full URL, which must start either with an http:// or https:// prefix. + To specify a URL that contains a question mark, press ESC followed by the "?". + If a https:// URL is specified, then ONTAP verifies the identity of the destination host by validating its certificate. If the Online Certificate Status Protocol (OCSP) is enabled for Event Management System (EMS), then ONTAP uses that protocol

to determine the certificate's revocation status. Use the `security config oscp show -application ems` command to determine if the OCSP-based certificate revocation status check is enabled for EMS. The parameter is specified only when the event notification destination type is already "rest-api". It is not allowed to specify the parameter for a destination that already has another type of destination address.

`[-certificate-authority <text>]` - Client Certificate Issuing CA

Use this parameter to specify a new value of the certificate authority (CA) to replace the current value in the event notification destination. There can be multiple client certificates installed for the admin vserver in the cluster, and this parameter, along with the *`certificate-serial`* parameter, uniquely identifies which one. + Use the security certificate show command to see the list of certificates installed in the cluster.

`[-certificate-serial <text>]` - Client Certificate Serial Number }

Use this parameter to specify a new serial number of the client certificate to replace the current value in the event notification destination.

`[-access-control-role <text>]` - Access Control Role (privilege: advanced)

Use this parameter to specify a new access control role to replace the current value in the event notification destination.

Examples

The following example shows the modification of event notification destinations:

```
cluster1::> event notification destination show

Name            Type        Destination
--------------  ----------  ---------------------
StorageAdminEmail
                email       Storage@example.com
StorageAdminSyslog
                syslog      example.com
snmp-traphost   snmp        10.30.40.10 (from "system snmp traphost")
3 entries were displayed.

cluster1::> event notification destination modify -name StorageAdminEmail
-email StorageAdmin@example.com

cluster1::> event notification destination show

Name            Type        Destination
--------------  ----------  ---------------------
StorageAdminEmail
                email       StorageAdmin@example.com
StorageAdminSyslog
                syslog      example.com
snmp-traphost   snmp        10.30.40.10 (from "system snmp traphost")
3 entries were displayed.
```

The following example shows how to clear the client certificate configuration when mutual authentication with the REST API server is no longer required:

```
cluster1::> event notification destination show -name RestApi -instance
Destination Name: RestApi
             Type of Destination: rest-api
              Destination Values: https://rest.example.com/rest
   Client Certificate Issuing CA: cluster1-root-ca
Client Certificate Serial Number: 052213E60B7088
       Client Certificate Valid?: -
                     Syslog Port: -
                Syslog Transport: -
           Syslog Message Format: -
Syslog Timestamp Format Override: -
 Syslog Hostname Format Override: -
      System-Defined Destination: false

cluster-1::> event notification destination modify -name RestApi
-certificate-authority - -certificate-serial -

cluster-1::> event notification destination show -name RestApi -instance
Destination Name: RestApi
             Type of Destination: rest-api
              Destination Values: https://rest.example.com/rest
   Client Certificate Issuing CA: -
Client Certificate Serial Number: -
       Client Certificate Valid?: -
                     Syslog Port: -
                Syslog Transport: -
           Syslog Message Format: -
Syslog Timestamp Format Override: -
 Syslog Hostname Format Override: -
      System-Defined Destination: false
```

Related Links

- event notification destination create
- security certificate show

event notification destination prepare-for-revert

Deletes or updates unsupported syslog destinations (transport=TCP or transport=UDP with non-default configurations: port, message-format, timestamp-format-override, hostname-format-override)

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

Description

The `event notification destination prepare-for-revert` can be used to remove or update syslog notification destinations that are not supported when the cluster reverts to the previous release. Supported syslog destinations are ones with *`udp-unencrypted`*`syslog-transport` using `syslog-port` *`514 and legacy_netapp`*`syslog-message-format` with `syslog-timestamp-format-override` and `syslog-hostname-format-override` both set to *`no-override`* . Syslog destinations with any other configurations are not supported.

Parameters

{ `-delete-unsupported-syslog-destinations {true|false}` - Clear unsupported syslog destinations (privilege: advanced)

Use this parameter to delete syslog destinations that are not supported in the previous release.

| `-update-unsupported-syslog-destinations {true|false}` - Update unsupported syslog destinations to supported (privilege: advanced) }

Use this parameter to update syslog destinations that are not supported in the previous release with supported configurations.

Examples

The following shows examples of "event notification destination prepare-for-revert":

```
cluster1::*> event notification destination show
Name            Type        Destination
--------------  ----------  ---------------------
snmp-traphost   snmp        - (from "system snmp traphost")
tst01           syslog      test.com (port: 6514, transport: tcp-
encrypted)
tst02           syslog      test.com (port: 601, transport: tcp-
unencrypted)
tst03           syslog      test.com (port: 1234, transport: udp-
unencrypted)
tst04           syslog      test.com (port: 514, transport: udp-
unencrypted)
5 entries were displayed.

cluster1::*> event notification destination prepare-for-revert -delete
-unsupported-syslog-destinations true

cluster1::*> event notification destination show
Name            Type        Destination
--------------  ----------  ---------------------
snmp-traphost   snmp        - (from "system snmp traphost")
tst04           syslog      test.com (port: 514, transport: udp-
unencrypted)
2 entries were displayed.
```

event notification destination show

Display event notification destinations

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

Description

The `event notification destination show` command displays event notification destinations. More detailed information about parameters can be found in the man page for the event notification destination create command.

Note: In the case of a rest-api destination type or syslog destination type (with tcp-encrypted transport), Online Certificate Status Protocol (OCSP) information is not included. OCSP information is available in the security config ocsp show -app ems command.

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.

[-name <text>] - Destination Name

Use this optional parameter to display information of an event notification destination that has the specified name.

[-type {snmp|email|syslog|rest-api}] - Type of Destination

Use this optional parameter to display information of event notification destinations that have the specified destination type.

[-destination <text>,…] - Destination

Use this optional parameter to display information of event notification destinations that have the specified destination address. Enter multiple addresses separated by a comma.

[-server-ca-present {true|false}] - Server CA Certificates Present?

Use this optional parameter to display information of event notification destinations that have the specified server-ca-present value. This field indicates whether there are certificates of the server-ca type exist in the system. If not, event messages will not be sent to a rest-api type destination having an HTTPS URL.

[-certificate-authority <text>] - Client Certificate Issuing CA

Use this optional parameter to display information of event notification destinations that have the specified certificate authority name.

[-certificate-serial <text>] - Client Certificate Serial Number

Use this optional parameter to display information of event notification destinations that have the specified certificate serial number.

[-certificate-valid {true|false}] - Client Certificate Valid?

Use this optional parameter to display information of event notification destinations that have the specified certificate-valid value. This field indicates whether the client certificate specified by the certificate-authority and certificate-serial fields is valid. If not, and if the REST API server requires client authentication, event messages are not sent to the server.

[-syslog-port <integer>] - Syslog Port

Use this optional parameter to display information about an event notification destination that has the specified syslog port.

[-syslog-transport {udp-unencrypted|tcp-unencrypted|tcp-encrypted}] - Syslog Transport

Use this optional parameter to display information about an event notification destination that has the specified syslog transport.

[-syslog-message-format {legacy-netapp|rfc-5424}] - Syslog Message Format

Use this optional parameter to display information about an event notification destination that has the specified syslog message format.

[-syslog-timestamp-format-override {no-override|rfc-3164|iso-8601-utc|iso-8601-local-time}] - Syslog Timestamp Format Override

Use this optional parameter to display information about an event notification destination that has the specified syslog timestamp format override.

[-syslog-hostname-format-override {no-override|fqdn|hostname-only}] - Syslog Hostname Format Override

Use this optional parameter to display information about an event notification destination that has the specified syslog hostname format override.

[-system-defined {true|false}] - System-Defined Destination

Use this optional parameter to display information about an event notification destination that has the specified system-defined value.

[-access-control-role <text>] - Access Control Role (privilege: advanced)

Use this optional parameter to display information about an event notification destination that has the specified access control role.

Examples

The following shows examples of "event notification destination show":

```
cluster1::> event notification destination show

Name            Type        Destination
--------------  ----------  ---------------------
StorageAdminEmail
                email       StorageAdmin@example.com (via "localhost" from
"admin@localhost", configured in "event config")
StorageAdminSyslog
                syslog      example.com (port: 514, transport: udp-
unencrypted)
snmp-traphost   snmp        10.30.40.10 (from "system snmp traphost")
RestApi         rest-api    https://rest.example.com/rest
4 entries were displayed.

cluster1::> event notification destination show -type snmp -instance
Destination Name: snmp-traphost
             Type of Destination: snmp
                     Destination: 10.30.40.10 (from "system snmp
traphost")
 Server CA Certificates Present?: -
   Client Certificate Issuing CA: -
Client Certificate Serial Number: -
       Client Certificate Valid?: -
                     Syslog Port: -
                Syslog Transport: -
           Syslog Message Format: -
Syslog Timestamp Format Override: -
 Syslog Hostname Format Override: -
      System-Defined Destination: false
```

Related Links

- event notification destination create
- security config ocsp show

event notification history show

Display latest events sent to destination

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

Description

The `event notification history show` command displays a list of event messages that have been sent to a notification destination. Information displayed by the command for each event is identical to that of the event log show command. This command displays events sent to a notification destination while the event log show command displays all events that have been logged.

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.

-destination <text> - Destination

Specifies the destination to which event messages have been sent to be displayed.

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

Displays a list of events for the node you specify. Use this parameter with the -seqnum parameter to display detailed information.

[-seqnum <Sequence Number>] - Sequence#

Selects the events that match this parameter value. Use with the -node parameter to display detailed information.

[-time <MM/DD/YYYY HH:MM:SS>] - Time

Selects the events that match this parameter value. Use the format: MM/DD/YYYY HH:MM:SS [+- HH:MM]. You can specify a time range by using the ".." operator between two time statements.

[-severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}] - Severity

Selects the events that match this parameter value. Severity levels are as follows:

- EMERGENCY - Disruption.
- ALERT - Single point of failure.
- ERROR - Degradation.
- NOTICE - Information.
- INFORMATIONAL - Information.
- DEBUG - Debug information.

[-message-name <Message Name>] - Message Name

Selects the events that match this parameter value (string). Message names are descriptive, so filtering output by message name displays messages of a specific type.

[-event <text>] - Event

Selects the events that match this parameter value. This parameter is useful when entered with wildcards. The "event" field contains the full text of the event, including any parameters. For example, the wafl.vol.offline event displays the name of the volume that is taken offline.

Examples

The following example displays all the events which match "important-events" filter and forwarded to the "snmp-traphost" destination:

```
cluster1::> event filter show
Filter      Rule Rule                                     SNMP Trap
Name        Posn Type     Message Name     Severity      Type
Parameters
----------- ---- -------- ---------------- ------------- ---------
-----------
default-trap-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  *                *             Standard, Built-
in
                                                                    *=*
            3    exclude  *                *             *         *=*
important-events
            1    include  *                EMERGENCY, ALERT
                                                          *         *=*
            2    include  callhome.*       ERROR         *         *=*
            3    exclude  *                *             *         *=*
no-info-debug-events
            1    include  *                EMERGENCY, ALERT, ERROR, NOTICE
                                                          *         *=*
            2    exclude  *                *             *         *=*
8 entries were displayed.

cluster1::> event notification destination show
Name            Type        Destination
--------------  ----------  ---------------------
snmp-traphost   snmp        192.168.10.40 (from "system snmp traphost")

cluster1::> event notification show
ID      Filter Name       Destinations
------- ----------------  -----------------
1       important-events  snmp-traphost

cluster1::>event notification history show -destination snmp-traphost
Time                Node             Severity      Event
------------------- ---------------- -------------
---------------------------
5/14/2015 03:02:09  node1            EMERGENCY     callhome.clam.node.ooq:
Call home for NODE(S) OUT OF CLUSTER QUORUM.
5/13/2015 12:05:45  node1            ALERT         od.rdb.mbox.read.error:
message="RDB-HA readPSlot: Failed to read blob_type 19, (pslot 16),
instance 1: 1 (1)."
2 entries were displayed.
```

event role-config commands

event role-config create

Create role-based event configuration

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

Description

The `event role-config create` command creates an EMS role-based configuration for an access control role. It provides the ability to assign an event filter to an access control role. Once an event filter is assigned to the access control role, only a limited subset of event management system (EMS) messages that match the event filter are visible to users of that role and only those limited subset of messages are sent as notifications to them. The assigned filter is applied transparently in both cases. The command also provides the ability to limit access to global EMS configurations available with the "event config" commands. Limiting access to EMS events and configurations is typically applied for an access control role that is designed to have limited administrative capabilities.

Parameters

`-access-control-role <text>` - Access Control Role (privilege: advanced)

Use this mandatory parameter to specify the access control role of the EMS role-based configuration.

`[-filter-name <text>]` - Event Filter Name (privilege: advanced)

Use this optional parameter to specify the name of the event filter that will be assigned to the access control role.

`[-limit-access-to-global-configs {true|false}]` - Limit Access to Global Configs (privilege: advanced)

Use this optional parameter to limit access to the global EMS configurations available with the "event config" commands. If no value is provided this field is set to true by default.

Examples

The following examples create role-based event configurations:

```
cluster1::> event role-config create -access-control-role storage-admin
             -filter-name storage-admin-events -limit-access-to-global
-configs true

cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true

cluster1::> event role-config create -access-control-role storage-admin
             -filter-name storage-admin-events

cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true

cluster1::> event role-config create -access-control-role storage-admin
             -limit-access-to-global-configs false

cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       -                     false
```

event role-config delete

Delete role-based event configuration

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

Description

The `event role-config delete` command deletes the EMS role-based configuration of an access control role.

Parameters

`-access-control-role <text>` - Access Control Role (privilege: advanced)

Use this mandatory parameter to specify the access control role for which the EMS role-based configuration needs to be deleted.

Examples

The following example shows the deletion of a role-based event configuration:

```
cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true

cluster1::> event role-config delete -access-control-role storage-admin

cluster1::> event role-config show
This table is currently empty.
```

event role-config modify

Modify role-based event configuration

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

Description

The `event role-config modify` command updates the EMS role-based configuration of an access control role.

Parameters

`-access-control-role <text>` - Access Control Role (privilege: advanced)

Use this mandatory parameter to specify the access control role for which the EMS role-based configuration needs to be modified.

`[-filter-name <text>]` - Event Filter Name (privilege: advanced)

Use this parameter to specify the new event filter name that needs to be assigned to the access control role.

`[-limit-access-to-global-configs {true|false}]` - Limit Access to Global Configs (privilege: advanced)

Use this parameter to change the limited access to global EMS configurations available with the "event config" commands.

Examples

The following examples show the modification of role-based event configurations:

```
cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true

cluster1::> event role-config modify -access-control-role storage-admin
             -filter-name storage-admin-events2

cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events2 true
cluster1::> event role-config modify -access-control-role storage-admin
             -filter-name storage-admin-events -limit-access-to-global
-configs false

cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  false
cluster1::> event role-config modify -access-control-role storage-admin
             -limit-access-to-global-configs true

cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true
```

event role-config show

Display the list of existing role-based event configurations

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

Description

The `event role-config show` command displays the EMS role-based configurations. It shows the list of access control roles with the event filters that are assigned to each role and the indication whether the access control role has limited access to global EMS configurations available with the "event config" commands.

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.

[-access-control-role <text>] - Access Control Role (privilege: advanced)

Use this parameter to only display the EMS role-based configurations assigned to this role.

[-filter-name <text>] - Event Filter Name (privilege: advanced)

Use this parameter to display all the access control roles that this filter is assigned to.

[-limit-access-to-global-configs {true|false}] - Limit Access to Global Configs (privilege: advanced)

Use this parameter to display all the access control roles that have this value for limited access to global EMS configurations.

Examples

The following example displays the role-based event configurations:

```
cluster1::> event role-config show
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true
storage-admin2      storage-admin-events  false
The following example displays the role-based event config for
a specific access control role:

cluster1::*> event role-config show -access-control-role storage-admin2
Access Control Role: storage-admin2
             Event Filter Name: storage-admin-events
Limit Access to Global Configs: false
The following example displays all the access control roles that
a specific filter is assigned to:

cluster1::*> event role-config show -filter-name storage-admin-events
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true
storage-admin2      storage-admin-events  false
2 entries were displayed.
The following example displays all the access control roles that have a
specific value for limited access to EMS global configurations:

cluster1::*> event role-config show -limit-access-to-global-configs true
Access Control Role Filter Name           Limit Access to Global Configs
------------------- --------------------- ------------------------------
storage-admin       storage-admin-events  true
```

event status commands

event status show

Display event status

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

Description

The `event status show` command summarizes information about occurrences of events. For detailed information about specific occurrences of events, use the event log show command.

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>|local}] - Node

Selects the event records that match this parameter value. Events are tracked on a node-by-node basis, rather than being rolled up cluster-wide.

[-message-name <Message Name>] - Message Name

Selects the event records that match this parameter value. The message name is a short descriptive string. Filtering output by message name displays messages of a specific type.

[-indications <integer>] - Number of Indications

Selects the event records that match this parameter value. This parameter is most useful when used with a range, such as using the range ">20" to display only events that have been posted more than 20 times.

[-drops <integer>] - Number of Drops

Selects the event records that match this parameter value.

[-last-time-occurred <MM/DD/YYYY HH:MM:SS>] - Last Indication Time

Selects the event records that match this parameter value.

[-last-time-dropped <MM/DD/YYYY HH:MM:SS>] - Last Suppressed Indication Time

Selects the event records that match this parameter value.

[-last-time-processed <MM/DD/YYYY HH:MM:SS>] - Last Processed Indication Time

Selects the event records that match this parameter value.

[-stat-starting-time <MM/DD/YYYY HH:MM:SS>] - Stat Starting Time

Selects the event records that match this parameter value.

[-last-hour-histogram <integer>,…] - 60-minute Histogram (privilege: advanced)

Use this parameter with the `-fields` parameter to display the "last hour" histogram for each event type. The last hour histogram records the number of times each event occurred in the last hour. The histogram is divided into sixty buckets, and each bucket collects one minute's events. The buckets display with the most recent event first.

[-last-day-histogram <integer>,…] - 24-hour Histogram (privilege: advanced)

Use this parameter with the `-fields` parameter to display the "last day" histogram for each event type. The last day histogram records the number of times each event occurred in the last day. The histogram is divided into 24 buckets, and each bucket collects one hour's events. The buckets display with the most recent event first.

[-last-week-histogram <integer>,…] - 7-day Histogram (privilege: advanced)

Use this parameter with the `-fields` parameter to display the "last week" histogram for each event type. The last week histogram records the number of times each event occurred in the last week. The histogram is divided into 7 buckets, and each bucket collects one day's events. The buckets display with the most recent event first.

[-severity {NODE_FAULT|SVC_FAULT|NODE_ERROR|SVC_ERROR|WARNING|NOTICE|INFO|DEBUG|VAR}] - Severity

Selects events that have the event severity you specify. Severity levels sort with the most severe levels first. Severity levels:

- NODE_FAULT - The node has detected data corruption, or is unable to provide client service.
- SVC_FAULT - The node has detected a temporary loss of service. Typically, this is caused by a transient software fault.
- NODE_ERROR - The node has detected a hardware error that is not immediately fatal.
- SVC_ERROR - The node has detected a software error that is not immediately fatal.
- WARNING - A high-priority message that does not indicate a fault.
- NOTICE - A normal-priority message that does not indicate a fault.
- INFO - A low-priority message that does not indicate a fault.
- DEBUG - A debugging message. These messages are typically suppressed.
- VAR - These messages have variable severity. Severity level for these messages is selected at runtime.

The examples below illustrate how to query on severity.

Examples

The following example displays recent event-occurrence status for node1:

```
cluster1::> event status show -node node1
Node               Message                       Occurs Drops Last Time
------------------ ----------------------------- ------ -----
-------------------
node1              raid.spares.media_scrub.start
                                                 6      0     3/11/2010
15:59:00
node1              raid.uninitialized.parity.vol
                                                 3      0     3/11/2010
15:58:28
node1              raid.vol.state.online         3      0     3/11/2010
15:58:29
node1              reg.defaultCommit.set.timeTaken
                                                 1      0     3/11/2010
15:58:28
node1              scsitgt.ha.state.changed      2      0     3/11/2010
15:58:28
node1              ses.multipath.notSupported    2      0     3/11/2010
15:58:43
node1              shelf.config.mpha             1      0     3/11/2010
15:58:48
node1              sk.hog.runtime                1      0     3/11/2010
15:58:28
node1              snmp.agent.msg.access.denied 1      0     3/11/2010
15:58:28
node1              snmp.link.up                  6      0     3/11/2010
15:58:28
node1              tar.csum.mismatch             2      0     3/11/2010
15:58:28
node1              tar.extract.success           2      0     3/11/2010
15:58:28
node1              vifmgr.lifsuccessfullymoved  3      0     3/11/2010
15:58:46
node1              vifmgr.portdown               1      0     3/11/2010
15:58:48
node1              vifmgr.portup                 5      0     3/11/2010
15:58:48
node1              vifmgr.startedsuccessfully    1      0     3/11/2010
15:58:43
```

The following example displays a summary of events which are warnings or more severe:

```
cluster1::> event status show -node node1 -severity <=warning -fields
indications,drops,severity
node    message-name             indications drops severity
------- ------------------------ ----------- ----- --------
node1   api.output.invalidSchema 5463        840   WARNING
node1   callhome.dsk.config      1           0     WARNING
node1   callhome.sys.config      1           0     SVC_ERROR
node1   cecc_log.dropped         145         0     WARNING
node1   cecc_log.entry           5           0     WARNING
node1   cecc_log.entry_no_syslog 4540        218   WARNING
node1   cecc_log.summary         5           0     WARNING
node1   cf.fm.noPartnerVariable  5469        839   WARNING
node1   cf.fm.notkoverBadMbox    1           0     WARNING
node1   cf.fm.notkoverClusterDisable 1       0     WARNING
node1   cf.fsm.backupMailboxError 1          0     WARNING
node1   cf.takeover.disabled     23          0     WARNING
node1   cmds.sysconf.logErr      1           0     NODE_ERROR
node1   config.noPartnerDisks    1           0     NODE_ERROR
node1   fci.initialization.failed 2          0     NODE_ERROR
node1   fcp.service.adapter      1           0     WARNING
node1   fmmb.BlobNotFound        1           0     WARNING
node1   ha.takeoverImpNotDef     1           0     WARNING
node1   httpd.config.mime.missing 2          0     WARNING
node1   mgr.opsmgr.autoreg.norec 1           0     WARNING
node1   monitor.globalStatus.critical 1      0     NODE_ERROR
node1   raid.mirror.vote.versionZero 1       0     SVC_ERROR
node1   ses.multipath.notSupported 2         0     NODE_ERROR
node1   snmp.agent.msg.access.denied 1       0     WARNING
24 entries were displayed.
```

The above example makes use of several features which are common to all `show` commands:

- A query is specified for the severity parameter. A query restricts the output of the show command; only rows matching the query will be displayed. In this case, the query indicates that only events which have a severity of "WARNING" or more severe will be displayed.
- The fields parameter selects the fields to display. Note that the severity field is not displayed in the default output.

Related Links

- event log show

job commands

job delete

Delete a job

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

Description

The `job delete` command deletes a job. Use the job show command to view a list of running jobs that can be deleted.

Parameters

`-id <integer>` - Job ID

The numeric ID of the job you want to delete. A job ID is a positive integer.

`[-vserver <vserver name>]` - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example deletes the job that has ID 99:

```
cluster1::> job delete -id 99
```

Related Links

- job show

job pause

Pause a job

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

Description

The `job pause` command pauses a job. Use the job resume command to resume a paused job. Use the job show command to view a list of running jobs that can be paused.

Parameters

`-id <integer>` - Job ID

The numeric ID of the job you want to pause. A job ID is a positive integer.

[-vserver <vserver name>] - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example pauses the job that has ID 183:

```
cluster1::> job pause -id 183
```

Related Links

- job resume
- job show

job resume

Resume a job

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

Description

The `job resume` command resumes a job that was previously paused by using the job pause command. Use the job show command to view a list of paused jobs that can be resumed.

Parameters

-id <integer> - Job ID

The numeric ID of the paused job to be resumed. A job ID is a positive integer.

[-vserver <vserver name>] - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example resumes the paused job that has ID 183:

```
cluster1::> job resume -id 183
```

Related Links

- job pause
- job show

job show-bynode

Display a list of jobs by node

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

Description

The `job show-bynode` command displays information about jobs on a per-node basis. The command output depends on the parameters specified with the command. If no parameters are specified, the command displays information about all jobs in the cluster that are currently owned by a node.

To display detailed information about a specific job, run the command with the `-id` parameter. The detailed view includes all of the default information plus additional items.

You can specify additional parameters to display only information that matches the values you specify for those parameters. For example, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

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>|local}] - Node

Use this parameter to display information only about the jobs that are associated with the node you specify.

[-id <integer>] - Job ID

Use this parameter to display information only about the jobs that match the ID or range of IDs you specify.

[-vserver <vserver name>] - Owning Vserver

Use this parameter with the name of a Vserver to display only jobs that are owned by that Vserver.

[-name <text>] - Name

Use this parameter to display information only about the jobs that match the job name you specify.

[-description <text>] - Description

Use this parameter to display information only about the jobs that match the description you specify.

[-affinity {Cluster|Node}] - Affinity

Use this parameter with an affinity value to display only jobs that match the affinity you specify.

[-username <text>] - User Name

Use this parameter with a username to display only jobs that are associated with that user.

Examples

The following example displays information about all jobs on a per-node basis:

```
node::> job show-bynode
                                        Owning
Node            Job ID Name             Vserver    Affinity
--------------- ------ --------------- ---------- --------
node0           1501   log-rotation    node-vserver
                                                   Cluster
                Descr:logrotation job
node1           1498   log-rotation    node-vserver
                                                   Cluster
                Descr:logrotation job
node2           1499   log-rotation    node-vserver
                                                   Cluster
                Descr:logrotation job
node3           1500   log-rotation    node-vserver
                                                   Cluster
                Descr:logrotation job
```

job show-cluster

Display a list of cluster jobs

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

Description

The `job show-cluster` command displays information about cluster-affiliated jobs. The command output depends on the parameters specified with the command. If no parameters are specified, the command displays information about all cluster-affiliated jobs.

To display detailed information about a specific job, run the command with the `-id` parameter. The detailed view includes all of the default information plus additional items.

You can specify additional parameters to display only information that matches the values you specify for those parameters. For example, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

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.

[-id <integer>] - Job ID

Use this parameter to display information only about the jobs that match the ID or range of IDs you specify.

[-vserver <vserver name>] - Owning Vserver

Use this parameter with the name of a Vserver to display only jobs that are owned by that Vserver.

[-name <text>] - Name

Use this parameter to display information only about the jobs that match the job name you specify.

[-description <text>] - Description

Use this parameter to display information only about the jobs that match the description you specify.

[-priority {Low|Medium|High|Exclusive}] - Priority

Use this parameter to display information only about the jobs that match the priority you specify.

[-node <nodename>] - Node

Use this parameter to display information only about the jobs that are associated with the node you specify.

[-affinity {Cluster|Node}] - Affinity

Use this parameter with an affinity value to display only jobs that match the affinity you specify.

[-schedule <job_schedule>] - Schedule

Use this parameter to display information only about the jobs that run on the schedule you specify.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time

Use this parameter to display information only about the jobs that match the queue time you specify.

[-starttime <MM/DD HH:MM:SS>] - Start Time

Use this parameter to display information only about the jobs that match the start time you specify.

[-endtime <MM/DD HH:MM:SS>] - End Time

Use this parameter to display information only about the jobs that match the end time you specify.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time

Use this parameter to display information only about the jobs that match the final timeout time you specify.

[-restarted {true|false}] - Restarted?

Use this parameter to display information only about the jobs that match the restart value you specify.

[-state {Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State

Use this parameter to display information only about the jobs that match the job state you specify.

[-code <integer>] - Status Code

Use this parameter to display information only about the jobs that match the status code you specify.

[-completion <text>] - Completion String

Use this parameter to display information only about the jobs that match the completion text you specify.

[-jobtype <text>] - Job Type

Use this parameter to display information only about the jobs that match the job type you specify.

[-category <text>] - Job Category

Use this parameter to display information only about the jobs that match the job category you specify.

[-uuid <UUID>] - UUID

Use this parameter to display information only about the jobs that match the UUID you specify.

[-username <text>] - User Name

Use this parameter with a username to display only jobs that are associated with the user you specify.

Examples

The following example displays information about all cluster-affiliated jobs:

```
cluster1::> job show-cluster
                        Owning
Job ID Name             Vserver   Node           State
------ ---------------- --------- -------------- -----------
305    Auto_Mirror      node-vserver
                                -              Running
6202   mirror-03_10     node-vserver
                                -              Queued
       Descr:Auto mirror
6203   mirror-04_10     node-vserver
                                -              Queued
       Descr:Auto mirror
6204   mirror-01_10     node-vserver
                                -              Queued
       Descr:Auto mirror
6205   mirror-02_10     node-vserver
                                -              Queued
       Descr:Auto mirror
6206   mirror-05_10     node-vserver
                                -              Queued
       Descr:Auto mirror
```

job show-completed

Display a list of completed jobs

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

Description

The `job show-completed` command displays information about completed jobs. The command output depends on the parameters you specify with the command. If you do not use any parameters, the command displays information about all completed jobs.

To display detailed information about a specific job, run the command with the `-id` parameter. The detailed view includes all of the default information plus additional items.

You can specify additional parameters to display only information that matches those parameters. For instance, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

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.

[-id <integer>] - Job ID

Use this parameter to display information only about the jobs that match the ID or range of IDs you specify.

[-vserver <vserver name>] - Owning Vserver

Use this parameter with the name of a Vserver to display only jobs that are owned by that Vserver.

[-name <text>] - Name

Use this parameter to display information only about the jobs that match the name you specify.

[-description <text>] - Description

Use this parameter to display information only about the jobs that match the description you specify.

[-priority {Low|Medium|High|Exclusive}] - Priority

Use this parameter to display information only about the jobs that match the priority you specify.

[-node <nodename>] - Node

Use this parameter to display information only about the jobs that are associated with the node you specify.

[-affinity {Cluster|Node}] - Affinity

Use this parameter with an affinity value to display only jobs that match the affinity you specify.

[-schedule <job_schedule>] - Schedule

If you use this parameter, the command displays information only about the jobs that have the schedule you specify.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time

If you use this parameter, the command displays information only about the jobs that have the queue time you specify.

[-starttime <MM/DD HH:MM:SS>] - Start Time

Use this parameter to display information only about the jobs that have the start time you specify.

[-endtime <MM/DD HH:MM:SS>] - End Time

Use this parameter to display information only about the jobs that have the end time you specify.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time

Use this parameter to display information only about the jobs that time out at the time you specify.

[-restarted {true|false}] - Restarted?

Use this parameter to display information only about the jobs that match the restart value you specify.

[-state {Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State

Use this parameter to display information only about the jobs that match the job state you specify.

[-code <integer>] - Status Code

Use this parameter to display information only about the jobs that match the status code you specify.

[-completion <text>] - Completion String

Use this parameter to display information only about the jobs that match the completion text you specify.

[-jobtype <text>] - Job Type

Use this parameter to display information only about the jobs that match the job type you specify.

[-category <text>] - Job Category

Use this parameter to display information only about the jobs that match the job category you specify.

[-uuid <UUID>] - UUID

Use this parameter to display information only about the jobs that match the UUID you specify.

[-username <text>] - User Name

Use this parameter with a username to display only jobs that are associated with that user.

Examples

The following example displays information about all completed jobs:

```
node::> job show-completed
                        Owning
Job ID Name             Vserver  End Time       Code       Completion
String
------ --------------- -------- -------------- ----------
--------------------
305     Auto_Mirror     node-vserver
                                 10/10 08:07:05 0          Succeeded
6202    mirror-03_10    node-vserver
                                 10/10 11:10:07 0
6203    mirror-04_10    node-vserver
                                 10/10 12:10:09 0
6204    mirror-01_10    node-vserver
                                 10/10 09:10:03 0
6205    mirror-02_10    node-vserver
                                 10/10 10:10:08 0
6206    mirror-05_10    node-vserver
                                 10/10 05:10:04 0
```

job show

Display a list of jobs

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

Description

The `job show` command displays information about jobs. By default, the command displays information about all current jobs.

To display detailed information about a specific job, run the command with the `-id` parameter.

You can specify additional parameters to select information that matches the values you specify for those parameters. For example, to display information only about jobs running on a specific node, run the command with the `-node` parameter.

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.

| [-inprogress]

Displays the job ID, the job name, the owning Vserver, and the progress of the job.

| [-jobstate]

Displays information about each job's state, including the queue state, whether the job was restarted, and when the job has completely timed out.

| [-jobuuid] (privilege: advanced)

Displays the job ID, the job name, the owning Vserver, and the job UUID.

| [-sched]

Displays the job ID, the job name, the owning Vserver, and the schedule on which the job runs.

| [-times]

Displays the job ID, the job name, the owning Vserver, the time when the job was last queued, the time when the job was last started, and the time when the job most recently ended.

| [-type]

Displays the job ID, the job name, the job type, and the job category.

| [-instance] }

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

[-id <integer>] - Job ID

Selects the jobs that match the ID or range of IDs that you specify.

[-vserver <vserver name>] - Owning Vserver

Selects jobs that are owned by the specified Vserver.

[-name <text>] - Name

Selects the jobs that match this parameter value.

[-description <text>] - Description

Selects the jobs that match this parameter value.

[-priority {Low|Medium|High|Exclusive}] - Priority

Selects the jobs that match this parameter value.

[-node <nodename>] - Node

Selects the jobs that match this parameter value.

[-affinity {Cluster|Node}] - Affinity

Selects the jobs that match this parameter value.

[-schedule <job_schedule>] - Schedule

Selects the jobs that match this parameter value.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time

Selects the jobs that match this parameter value.

[-starttime <MM/DD HH:MM:SS>] - Start Time

Selects the jobs that match this parameter value.

[-endtime <MM/DD HH:MM:SS>] - End Time

Selects the jobs that match this parameter value.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time

Selects the jobs that match this parameter value.

[-restarted {true|false}] - Restarted?

Selects the jobs that match this parameter value.

[-state {Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State

Selects the jobs that match this parameter value.

[-code <integer>] - Status Code

Selects the jobs that match this parameter value.

[-completion <text>] - Completion String

Selects the jobs that match this parameter value.

[-jobtype <text>] - Job Type

Selects the jobs that match this parameter value.

[-category <text>] - Job Category

Selects the jobs that match this parameter value.

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

Selects the jobs that match this parameter value.

[-progress <text>] - Execution Progress

Selects the jobs that match this parameter value.

[-username <text>] - User Name

Selects the jobs that match this parameter value.

[-restart-is-delayed-by-module <text>] - Restart Is Delayed by Module

Selects jobs which are or were delayed by the specified module during the restart. For example: MCC_SWITCHBACK

Examples

The following example displays information about all jobs on the node named node1:

```
cluster1::> job show -node node1
                        Owning
Job ID Name             Vserver   Node         State
------ ---------------- --------- ------------ ----------
308114 mirror-daily-3587206
                        node-vserver
                                  node1        Running
       Descr:Auto-replicate to 1 mirror(s)
308115 mirror-daily-3618985
                        node-vserver
                                  node1        Running
       Descr:Auto-replicate to 1 mirror(s)
308116 mirror-daily-3619010
                        node-vserver
                                  node1        Queued
       Descr:Auto-replicate to 1 mirror(s)
308117 mirror-daily-3749547
                        node-vserver
                                  node1        Queued
       Descr:Auto-replicate to 1 mirror(s)
4 entries were displayed.
```

job stop

Stop a job

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

Description

The `job stop` command stops a running job. A stopped job cannot be resumed. Use the job pause command to pause a job so that you can later resume it. Use the job show command to view a list of running jobs.

Parameters

`-id <integer>` - Job ID

The numeric ID of the job to stop. A job ID is a positive integer.

`[-vserver <vserver name>]` - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example stops the job that has ID 101:

```
cluster1::> job stop -id 101
```

Related Links

- job pause
- job show

job unclaim

Unclaim a cluster job

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

Description

The `job unclaim` command causes a cluster-affiliated job that is owned by an unavailable node to be unclaimed by that node. Another node in the cluster can then take ownership of the job. Use the job show-cluster command to obtain a list of cluster-affiliated jobs.

Parameters

`-id <integer>` - Job ID (privilege: advanced)

Use this parameter to specify the ID number of the job to unclaim.

`[-vserver <vserver name>]` - Owning Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example shows how to unclaim the cluster-affiliated job with the ID `27` that is owned by the Vserver `vs1` :

```
cluster1::*> job unclaim -vserver vs1 -id 27
```

Related Links

- job show-cluster

job watch-progress

Watch the progress of a job

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

Description

The `job watch-progress` command displays the progress of a job, and periodically updates that display. You can specify the frequency of the updates.

Parameters

`-id <integer>` - Job ID

Use this parameter to specify the numeric ID of the job to monitor.

`[-vserver <vserver name>]` - Owning Vserver

Use this parameter to specify the name of the Vserver that owns the job.

`[-interval <integer>]` - Refresh Interval (seconds)

Use this parameter to specify the number of seconds between updates.

Examples

The following example show how to monitor the progress of the job that has ID `222` on Vserver `vs0` . The progress display updates every 3 seconds.

```
cluster1::> job watch-progress -vserver vs0 -id 222 -interval 3
```

job history commands

job history show

Display a history of jobs

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

Description

The `job history show` command displays a history of completed jobs with newer entries displayed first. You can specify optional parameters to select information about job history items that match only those parameters. For example, to display information about jobs that were completed on February 27 at noon, run the command with `-endtime "02/27 12:00:00"` .

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>|local}] - Node

Selects the completed jobs that match this parameter value.

[-record <Sequence Number>] - Record ID

Selects the completed jobs that match the record ID or range of record IDs you specify. Note that record IDs are unique for each node, not for the cluster as a whole. As a result, there can be two records with the same record ID within the cluster.

[-vserver <vserver name>] - Owning Vserver

Selects the completed jobs that are owned by the Vserver you specify.

[-id <integer>] - Job ID

Selects the completed jobs that match this parameter value.

[-endtime <MM/DD HH:MM:SS>] - End Time

Selects jobs that completed at the time you specify. This parameter is most useful when used with a range of times.

[-starttime <MM/DD HH:MM:SS>] - Start Time

Selects completed jobs that were started at the time you specify. This parameter is most useful when used with a range of times.

[-name <text>] - Name

Selects the completed jobs that match this parameter value.

[-description <text>] - Description

Selects the completed jobs that match this parameter value.

[-code <integer>] - Status Code

Selects the completed jobs that match this parameter value. Each job defines its own status codes. The completion text is more informative, but support technicians may request this numeric code.

[-progress <text>] - Progress String

Selects the completed jobs that match this parameter value.

[-completion <text>] - Completion String

Selects the completed jobs that match this parameter value.

[-jobuuid <UUID>] - Job UUID (privilege: advanced)

Selects the completed jobs that match this parameter value.

[-event-type {Idle|Running|Succeeded|Failed|Paused|Stopped|Deleted|Error}] - Event Type

Selects the completed jobs that match this parameter value.

[-event-time <MM/DD HH:MM:SS>] - Event Time

Selects the completed jobs that match this parameter value. This parameter is most useful when used with a range of times.

[-error-code <integer>] - Job Manager Error Code

Selects the completed jobs that match this parameter value.

[-error-text <text>] - Job Manager Error Text

Selects the completed jobs that match this parameter value.

[-username <text>] - User Name

Selects the completed jobs that match this parameter value.

Examples

The following example displays information about all completed jobs:

```
cluster1::> job history show
                                Owning
Time           Node             Vserver    Name                 Event
Job ID
-------------- -------------- ---------- -------------------- ---------
------
08/23 08:58:24 node1          node1-vs   Vol Create           Succeeded
76
     Description: Create testvol
      Completion: Successful
08/23 08:58:22 node1          node1-vs   Vol Create           Running
76
     Description: Create testvol
08/22 08:16:36 node1          node1-vs   CLUSTER BACKUP AUTO weekly
                                                              Succeeded
4
     Description: Cluster Backup Job
08/22 08:15:49 node1          node1-vs   CLUSTER BACKUP AUTO weekly
                                                              Running
4
     Description: Cluster Backup Job
08/22 08:15:08 node1          node1-vs   CLUSTER BACKUP AUTO weekly
                                                              Idle
4
     Description: Cluster Backup Job
08/22 08:15:03 node1          node1-vs   CLUSTER BACKUP AUTO weekly
                                                              Running
4
     Description: Cluster Backup Job
6 entries were displayed.
```

The following example shows how to use a range with the "endtime" parameter to select only the events that ended between 8:15 and 8:16 on August 22nd.

```
cluster1::> job history show -endtime "08/22 08:15:00".."08/22 08:16:00"
                              Owning
Time           Node           Vserver    Name                 Event
Job ID
-------------- -------------- ---------- -------------------- ---------
------
08/22 08:15:49 node1          node1-vs   CLUSTER BACKUP AUTO weekly
                                                              Running
4
     Description: Cluster Backup Job
08/22 08:15:08 node1          node1-vs   CLUSTER BACKUP AUTO weekly
                                                              Idle
4
     Description: Cluster Backup Job
08/22 08:15:03 node1          node1-vs   CLUSTER BACKUP AUTO weekly
                                                              Running
4
     Description: Cluster Backup Job
3 entries were displayed.
```

job initstate commands

job initstate show

Display init state for job managers

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

Description

The `job initstate show` command displays information about the initialization states of job-manager processes.

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>|local}]` - Node (privilege: advanced)

Selects the nodes that match this parameter value.

[-process <process_name>] - Process Name (privilege: advanced)

Selects the nodes that match this parameter value.

[-initialized {true|false}] - Initialized? (privilege: advanced)

Selects the nodes that match this parameter value (`true` means initialized; `false` means not initialized).

[-cache-root <text>] - Cache Root (privilege: advanced)

Selects the nodes that match this parameter value.

[-siteid <UUID>] - Site ID (privilege: advanced)

Selects the nodes that match this parameter value.

[-hp-threads <integer>] - High Priority Threads (privilege: advanced)

Selects the nodes that have the number of high-priority threads you specify.

[-mp-threads <integer>] - Medium Priority Threads (privilege: advanced)

Selects the nodes that have the number of medium-priority threads you specify.

[-lp-threads <integer>] - Low Priority Threads (privilege: advanced)

Selects the nodes that have the number of low-priority threads you specify.

[-tx-interval <integer>] - Transaction Interval (privilege: advanced)

Selects the nodes that have the number of seconds you specify as their transaction interval.

[-initmsg <text>] - Initialization Message (privilege: advanced)

Selects the nodes that match this parameter value.

[-thread-initmsg <text>] - Thread Initialization Message (privilege: advanced)

Selects the nodes that match this parameter value. The thread initialization message contains information about thread status. If there is no information to communicate, this message is empty.

[-recovery-enabled {true|false}] - Job Failover Enabled? (privilege: advanced)

Selects the nodes that match this parameter value (`true` means enabled, `false` means not enabled).

[-ex-threads <integer>] - Exclusive Priority Threads (privilege: advanced)

Selects the nodes that match this parameter value.

Examples

The following example shows how to display general job-manager initialization-state information for a cluster.

```
cluster1::*> job initstate show
                                     HP   MP   LP   EX
Node              Process     Init? Thr  Thr  Thr  Thr  TX Int Failover?
----------------- ----------- ----- ---- ---- ---- ---- ------ ---------
node1             mgwd        true  2    3    5    8    300    true
node2             mgwd        true  2    3    5    8    300    true
2 entries were displayed.
```

The following example shows how to display detailed job-manager initialization-state information for a node named `node0` .

```
cluster1::*> job initstate show -instance -node node0
                      Node: node0
              Process Name: mgwd
              Initialized?: true
                Cache Root: /mroot/jm_cache
                   Site ID: 824e8f7d-f49-1d9-84af-00423b7352
     High Priority Threads: 2
   Medium Priority Threads: 3
      Low Priority Threads: 5
      Transaction Interval: 300
    Initialization Message: Initialized
      Are Threads Running?: -
     Job Failover Enabled?: true
Exclusive Priority Threads: 8
```

job private commands

job private delete

Delete a job

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

Description

The `job private delete` command deletes a private job. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

If you use this command on a job that does not support the delete operation, the command returns an error message.

Use the job private show command to view a list of private jobs that can be deleted.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the node with which the private job is associated.

`-id <integer>` - Job ID (privilege: advanced)

Use this parameter to specify the numeric ID of the private job to be deleted. A job ID is a positive integer.

`[-vserver <vserver name>]` - Owning Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example shows how to delete the job that has ID `273` from the node named `node2` :

```
cluster1::*> job private delete -node node2 -id 273
```

Related Links

- job private show

job private pause

Pause a job

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

Description

The `job private pause` command pauses a private job. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

If you use this command to pause a job that does not support it, the command returns an error message.

Use the job private resume command to resume a paused private job.

Use the job private show command to view a list of private jobs.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the node with which the private job is associated.

`-id <integer>` - Job ID (privilege: advanced)

Use this parameter to specify the numeric ID of the paused private job to be paused. A job ID is a positive integer.

`[-vserver <vserver name>]` - Owning Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example pauses the private job that has ID `99` on the node `node1` :

```
cluster1::*> jobs private pause -node node1 -id 99
```

Related Links

- job private resume
- job private show

job private resume

Resume a job

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

Description

The `job private resume` command resumes a private job that was paused by using the job private pause command. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

Use the job private show command to view a list of paused private jobs that can be resumed.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the node with which the paused private job is associated.

`-id <integer>` - Job ID (privilege: advanced)

Use this parameter to specify the numeric ID of the paused private job to be resumed. A job ID is a positive integer.

`[-vserver <vserver name>]` - Owning Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example resumes the paused private job that has ID `99` on a node named `node2` :

```
cluster1::*> job private resume -node node2 -id 99
```

Related Links

- job private pause
- job private show

job private show-completed

Display a list of completed jobs

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

Description

The `job private show-completed` command displays information about completed private jobs. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

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>|local}] - Node (privilege: advanced)

Use this parameter to display information only about completed jobs that are associated with the node you specify.

[-id <integer>] - Job ID (privilege: advanced)

Use this parameter to display information only about completed jobs that have the ID you specify.

[-vserver <vserver name>] - Owning Vserver (privilege: advanced)

Use this parameter to display only completed jobs that are owned by the Vserver you specify.

[-name <text>] - Name (privilege: advanced)

Use this parameter to display information only about completed jobs that have the name you specify.

[-description <text>] - Description (privilege: advanced)

Use this parameter to display information only about completed jobs that have the description you specify.

[-priority {Low|Medium|High|Exclusive}] - Priority (privilege: advanced)

Use this parameter to display information only about completed jobs that have the priority you specify.

[-schedule <job_schedule>] - Schedule (privilege: advanced)

Use this parameter to display information only about completed jobs that have the schedule you specify.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the queue time you specify.

[-starttime <MM/DD HH:MM:SS>] - Start Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the start time you specify.

[-endtime <MM/DD HH:MM:SS>] - End Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the end time you specify.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time (privilege: advanced)

Use this parameter to display information only about completed jobs that have the final timeout time you specify.

[-restarted {true|false}] - Restarted? (privilege: advanced)

Use this parameter to display information only about completed jobs that have the restart value you specify.

[-state {Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State (privilege: advanced)

Use this parameter to display information only about completed jobs that have the job state you specify.

[-code <integer>] - Status Code (privilege: advanced)

Use this parameter to display information only about completed jobs that have the status code you specify.

[-completion <text>] - Completion String (privilege: advanced)

Use this parameter to display information only about completed jobs that have the completion text you specify.

[-jobtype <text>] - Job Type (privilege: advanced)

Use this parameter to display information only about completed jobs that have the job type you specify.

[-category <text>] - Job Category (privilege: advanced)

Use this parameter to display information only about completed jobs that have the job category you specify.

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

Use this parameter to display information only about completed jobs that have the UUID you specify.

[-username <text>] - User Name (privilege: advanced)

Use this parameter to display information only about completed jobs that are associated with the user you specify.

Examples

The following example shows how to display information about all completed private jobs on the node named `node1`:

```
cluster1::*> job private show-completed -node node1
Node: node1
                      Owning
Job ID Name           Vserver    End Time       Code     Completion String
------ -------------- ---------- -------------- --------
---------------------
1      sync task      node1      02/17 15:03:23 0
2      load_balancing node1      02/17 16:29:28 0        DONE_VIF_STATS
3      snap-hourly    node1      02/17 16:05:00 0
4      snap-daily     node1      02/17 00:10:00 0
5      snap-weekly    node1      02/13 00:15:00 0
8      Cross-Cluster Manager node1 02/17 16:27:27 0      complete
9      reconcile service policy node1 02/17 15:03:12 0
7 entries were displayed.
```

job private show

Display a list of jobs

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

Description

The `job private show` command displays information about private jobs. Private jobs are affiliated with a specific node and do not use any cluster facilities, such as the replicated database.

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.

| [-inprogress] (privilege: advanced)

Displays the job ID, name, owning Vserver, and progress of each private job.

| [-jobstate] (privilege: advanced)

Displays information about each private job's state, including the queue state, whether the job was restarted, and when the job has timed out.

| [-jobuuid] (privilege: advanced)

Displays the ID, name, owning Vserver, and UUID of each private job.

| [-sched] (privilege: advanced)

Displays the job ID, name, owning Vserver, and run schedule of each private job.

| [-times] (privilege: advanced)

Displays the queue time, start time, and end time of each private job.

| [-type] (privilege: advanced)

Displays the type and category of each private job.

| [-instance] (privilege: advanced) }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the private jobs that match this parameter value. .

[-id <integer>] - Job ID (privilege: advanced)

Selects the private jobs that match the ID or range of IDs that you specify.

[-vserver <vserver name>] - Owning Vserver (privilege: advanced)

Selects the private jobs that match this parameter value.

[-name <text>] - Name (privilege: advanced)

Selects the private jobs that match this parameter value.

[-description <text>] - Description (privilege: advanced)

Selects the private jobs that match this parameter value.

[-priority {Low|Medium|High|Exclusive}] - Priority (privilege: advanced)

Selects the private jobs that match this parameter value.

[-schedule <job_schedule>] - Schedule (privilege: advanced)

Selects the private jobs that match this parameter value.

[-queuetime <MM/DD HH:MM:SS>] - Queue Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-starttime <MM/DD HH:MM:SS>] - Start Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-endtime <MM/DD HH:MM:SS>] - End Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-dropdeadtime <MM/DD HH:MM:SS>] - Drop-dead Time (privilege: advanced)

Selects the private jobs that match this parameter value.

[-restarted {true|false}] - Restarted? (privilege: advanced)

Selects the private jobs that match this parameter value.

[-state {Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - State (privilege: advanced)

Selects the private jobs that match this parameter value.

[-code <integer>] - Status Code (privilege: advanced)

Selects the private jobs that match this parameter value.

[-completion <text>] - Completion String (privilege: advanced)

Selects the private jobs that match this parameter value.

[-jobtype <text>] - Job Type (privilege: advanced)

Selects the private jobs that match this parameter value.

[-category <text>] - Job Category (privilege: advanced)

Selects the private jobs that match this parameter value.

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

Selects the private jobs that match this parameter value.

[-progress <text>] - Execution Progress (privilege: advanced)

Selects the private jobs that match this parameter value.

[-username <text>] - User Name (privilege: advanced)

Selects the private jobs that match this parameter value.

Examples

The following example displays information about all private jobs on the local node:

```
cluster1::*> job private show -node local
Node: node1
                            Owning
Job ID Name                 Vserver    State
------ -------------------- ---------- -----------
3      snap-hourly          cluster1   Queued
       Description: Auto-Snapshot
4      snap-daily           cluster1   Queued
       Description: Auto-Snapshot
5      snap-weekly          cluster1   Queued
       Description: Auto-Snapshot
6      sync task            cluster1   Queued
       Description: sync task
7      ldap-certs           cluster1   Queued
       Description: ldap resync
5 entries were displayed.
```

job private stop

Stop a job

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

Description

The `job private stop` command stops a running private job. A private job is a job that is associated with a specific node and does not use cluster facilities. A stopped job cannot be restarted.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

This specifies the node on which the job is running.

`-id <integer>` - Job ID (privilege: advanced)

This specifies the numeric ID of the job that is to be stopped.

`[-vserver <vserver name>]` - Owning Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver that owns the job.

Examples

The following example stops a private job with the ID 416 on a node named node0:

```
cluster1::*> job private stop -node node0 -id 416
```

job private watch-progress

Watch the progress of a job

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

Description

The `job private watch-progress` command displays and periodically updates the progress of a private job. A private job is a job that is associated with a specific node and does not use cluster facilities. You can specify the frequency of the progress updates.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

This specifies the node on which the job is running.

`-id <integer>` - Job ID (privilege: advanced)

This specifies the numeric ID of the job whose progress is to be monitored.

`[-vserver <vserver name>]` - Owning Vserver (privilege: advanced)

Use this parameter to specify the Vserver with which the paused private job is associated.

`[-interval <integer>]` - Refresh Interval (seconds) (privilege: advanced)

This optionally specifies, in seconds, the frequency of the updates.

Examples

The following example monitors the progress of the private job that has ID 127 on a node named node1. The progress is updated every 2 seconds.

```
cluster1::*> job private watch-progress -node node1 -id 127 -interval 2
Queued
```

job schedule commands

job schedule delete

Delete a schedule

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

Description

The `job schedule delete` command deletes a schedule. Use the job schedule show command to display all current schedules.

You cannot delete any schedules that are in use by jobs. Use the job schedule show-jobs command to display jobs by schedule.

You cannot delete any schedules that are referenced by:

- Volume snapshot policy entries
- SnapMirror entries
- SIS policy entries
- configuration backup settings

You must remove all references to a schedule before you can delete it. If you attempt to delete a schedule that is referenced, an error message will list which entries reference the schedule you want to delete. Use the `show` command for each of the items listed by the error message to display which entries reference the schedule. You may need to use the `-instance` parameter to display more detail.

Parameters

`[-cluster <Cluster name>]` - Cluster

This parameter specifies the name of the cluster on which you want to delete a schedule. By default, the schedule is deleted from the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

`[-vserver <vserver name>]` - Vserver

This parameter specifies the name of the Vserver on which you want to delete a schedule.

`-name <text>` - Schedule Name

Use this parameter with the name of an existing schedule to specify the schedule you want to delete.

Examples

The following example deletes a schedule named overnightbackup:

```
cluster1::> job schedule delete -name overnightbackup
```

Related Links

- job schedule show
- job schedule show-jobs

job schedule show-jobs

Display the list of jobs by schedule

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

Description

The `job schedule show-jobs` command displays information about jobs that are associated with schedules.

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.

[-name <text>] - Schedule Name

Use this parameter to display information only about the jobs that are associated with the schedule you specify.

[-affinity {Cluster|Node}] - Cluster / Node

Use this parameter to display information only about the jobs that match the affinity value you specify.

[-owner <text>] - Owner

Use this parameter to display information only about the jobs that are owned by the nodes you specify.

[-jobid <integer>] - ID

Use this parameter to display information only about the jobs that match the ID or range of IDs that you specify.

[-jobname <text>] - Job Name

Use this parameter to display information only about the jobs that match the name you specify.

[-scheduleuuid <UUID>] - Schedule Uuid

displays the Uuid of the specific job schedule.

Examples

The following example shows information about schedules that are associated with jobs:

```
cluster1::> job schedule show-jobs
Name          Type      Owner                   Job ID     Job Name
Schedule Uuid
------------  --------- ---------------------  ---------- --------------
------------------
hourly        Cluster   -                       98644      mirror-hourly
8bafba5a-ff9a-11eb-8531-005056a75903
weeklylog     Node      node0                   1501       log-rotation
449c070c-ff9a-11eb-8531-005056a75903
weeklylog     Node      node1                   1498       log-rotation
8bb0adca-ff9a-11eb-8531-005056a75903
weeklylog     Node      node2                   1499       log-rotation
b15fce61-ff9a-11eb-8531-005056a75903
weeklylog     Node      node3                   1500       log-rotation
8bb14bd2-ff9a-11eb-8531-005056a75903
5 entries were displayed.
```

job schedule show

Display a list of available schedules

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

Description

The `job schedule show` command displays information about schedules.

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.

[-cluster <Cluster name>] - Cluster

Selects the schedules that match this parameter value.

[-vserver <vserver name>] - Vserver

Selects the schedules that match this parameter value. These job schedules can only be used by consumers in that Vserver.

[-name <text>] - Schedule Name

Selects the schedules that match this parameter value.

[-type {cron|interval|builtin}] - Schedule Type

Selects the schedules that match this parameter value.

[-description <text>] - Description

Selects the schedules that match this parameter value.

Examples

The following example displays information about all schedules:

```
cluster1::> job schedule show
Cluster       Vserver     Name         Type     Description
------------- ----------- ----------- ---------
-------------------------------------------
cluster1
              data_vs_1
                          5min        cron
@:00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55
                          daily       cron      @0:10
                          hourly      cron      @:05
                          monthly     cron      1@0:20
                          weekly      cron      Sun@0:15
```

job schedule cron create

Create a cron schedule

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

Description

The `job schedule cron create` command creates a cron schedule. A cron schedule, like a UNIX cron job, runs at a specified time. You can also specify months, days of the month, or days of the week on which the schedule will run.

If you specify values for both days of the month and days of the week, they are considered independently. For example, a cron schedule with the day specification Friday, 13 runs every Friday and on the 13th day of each month, not just on every Friday the 13th.

Parameters

[-cluster <Cluster name>] - Cluster

This parameter specifies the name of the cluster on which you want to create a cron schedule. By default, the schedule is created on the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

[-vserver <vserver name>] - Vserver

This parameter specifies the name of the Vserver on which you want to create a cron schedule.

-name <text> - Name

Use this parameter to specify the name of the cron schedule that you want to create.

[-month <cron_month>,…] - Month

Use this parameter to specify months in which the schedule runs. Valid values are January, February, March, April, May, June, July, August, September, October, November, December, and all. Specify "all" to run the schedule every month.

[-dayofweek <cron_dayofweek>,…] - Day of Week

Use this parameter to specify days of the week on which the schedule runs. Valid values are Sunday, Monday, Tuesday, Thursday, Friday, and Saturday, and all. Specify "all" to run the schedule every day.

[-day <cron_dayofmonth>,…] - Day

Use this parameter to specify days of the month on which the schedule runs. Valid values range from 1 to 31.

[-hour <cron_hour>,…] - Hour

Use this parameter to specify the hours value of the time of day at which the schedule runs. Valid values range from 0 (midnight) to 23 (11:00 p.m.). Specify "all" to run the schedule every hour.

-minute <cron_minute>,… - Minute

Use this parameter to specify the minutes portion of the time of day at which the schedule runs. Valid values range from 0 to 59.

Examples

The following example creates a cron schedule named weekendcron that runs on weekend days (Saturday and Sunday) at 3:00 a.m.

```
cluster1::> job schedule cron create -name weekendcron -dayofweek
"Saturday, Sunday" -hour 3 -minute 0
```

job schedule cron delete

Delete a cron schedule

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

Description

The `job schedule cron delete` command deletes a cron schedule. Use the job schedule cron show command to display all current cron schedules.

You cannot delete any cron schedules that are associated with jobs. Use the job schedule show-jobs command to display jobs by schedule.

Parameters

`[-cluster <Cluster name>]` - Cluster

This parameter specifies the name of the cluster on which you want to delete a cron schedule. By default, the schedule is deleted from the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

`[-vserver <vserver name>]` - Vserver

This parameter specifies the name of the Vserver on which you want to delete a cron schedule.

`-name <text>` - Name

Use this parameter with the name of an existing cron schedule to specify the cron schedule that you want to delete.

Examples

The following example deletes a cron schedule named midnightcron:

```
cluster1::> job schedule cron delete -name midnightcron
```

Related Links

- job schedule cron show
- job schedule show-jobs

job schedule cron modify

Modify a cron schedule

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

Description

The `job schedule cron modify` command modifies a cron schedule. A cron schedule, like a UNIX cron job, runs at a specified time. You can also specify months, days of the month, or days of the week on which the schedule runs. Use the job schedule cron show command to display all current cron schedules. See the documentation for job schedule cron show for more information about how cron schedules work.

Modifying one parameter of a cron schedule does not affect the other parameters. For example, if cron schedule is set to run at 3:15 AM, and you modify the "hour" parameter to 4, the schedule's new time will be 4:15am. To clear a parameter of the schedule's interval, you must explicitly set that portion to "0" or "-" Some parameters can also be set to "all".

Parameters

[-cluster <Cluster name>] - Cluster

Use this parameter to specify the cluster of an existing cron schedule you want to modify. The local cluster is provided as the default value. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

[-vserver <vserver name>] - Vserver

Use this parameter to specify the Vserver of an existing cron schedule you want to modify.

-name <text> - Name

Use this parameter with the name of an existing cron schedule to specify the cron schedule you want to modify.

[-month <cron_month>,…] - Month

Use this parameter to specify a new "month" value for the cron schedule. Valid values are January, February, March, April, May, June, July, August, September, October, November, December, or all. Specify "all" to run the schedule every month.

[-dayofweek <cron_dayofweek>,…] - Day of Week

Use this parameter to specify a new "day of week" value for the cron schedule. Valid values include Sunday, Monday, Tuesday, Thursday, Friday, Saturday, or all. Specify "all" to run the schedule every day.

[-day <cron_dayofmonth>,…] - Day

Use this parameter to specify a new "day of month" value for the cron schedule. Valid values range from 1 to 31.

[-hour <cron_hour>,…] - Hour

Use this parameter to specify a new "hour of the day" value for the cron schedule. Valid values range from 0 (midnight) to 23 (11:00 p.m.), Specify "all" to run the schedule every hour.

[-minute <cron_minute>,…] - Minute

Use this parameter to specify a new "minute of the hour" value for the cron schedule. Valid values range from 0 to 59.

Examples

The following example modifies a cron schedule named weekendcron so that it runs at 3:15 a.m.:

```
cluster1::> job schedule cron modify -name weekendcron -hour 3 -minute 15
```

Related Links

- job schedule cron show

job schedule cron show

Show cron schedules

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

Description

The `job schedule cron show` command displays information about cron schedules. A cron schedule runs a job at a specified time on specified days.

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.

[-cluster <Cluster name>] - Cluster

Selects the cron schedules that match this parameter value.

[-vserver <vserver name>] - Vserver

Selects the cron schedules that match this parameter value. These job schedules can only be used by consumers in that Vserver.

[-name <text>] - Name

Selects the cron schedules that match this parameter value.

[-month <cron_month>,…] - Month

Selects the cron schedules that match this parameter value. Valid values are `January` , `February` , `March` , `April` , `May` , `June` , `July` , `August` , `September` , `October` , `November` , `December` , or `all` .

[-dayofweek <cron_dayofweek>,…] - Day of Week

Selects the cron schedules that match this parameter value. Valid values include `Sunday` , `Monday` , `Tuesday` , `Wednesday` , `Thursday` , `Friday` , `Saturday` , or `all` .

[-day <cron_dayofmonth>,…] - Day

Selects the cron schedules that match this parameter value. Valid values range from `1` to `31` .

[-hour <cron_hour>,…] - Hour

Selects the cron schedules that match this parameter value.

[-minute <cron_minute>,…] - Minute

Selects the cron schedules that match the minute or range of minutes that you specify.

[-description <text>] - Description

Selects the cron schedules that match this parameter value.

Examples

The following example displays information about all current cron schedules:

```
cluster1::> job schedule cron show
Cluster       Vserver     Name        Description
------------- ----------- -----------
-----------------------------------------------------
cluster1
              data_vs_1
                          5min
@:00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55
                          8hour       @2:15,10:15,18:15
                          weekly      Sun@0:15
```

The following example displays information about the cron schedule named weekly:

```
cluster1::> job schedule cron show -name weekly -instance
Cluster: cluster1
    Vserver: data_vs_1
       Name: weekly
      Month: -
Day of Week: Sunday
        Day: -
       Hour: 0
     Minute: 15
Description: Sun@0:15
```

job schedule interval create

Create a schedule that runs on an interval

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

Description

The `job schedule interval create` creates an interval schedule. An interval schedule runs jobs at specified intervals after the previous job finishes. For instance, if a job uses an interval schedule of 12 hours and takes 30 minutes to complete, the job runs at the following times:

- Day one at 8:00 a.m. (the job's initial run)
- Day one at 8:30 p.m.
- Day two at 9:00 a.m.
- Day two at 9:30 p.m.

Each of the numerical parameters of the interval must be a whole number. These parameters can be used individually, or combined to define complex time values. For example, use a value of 1 day, 12 hours to create an interval of 1.5 days.

Large parameter values are converted into larger units. For example, if you create a schedule with an interval

of 36 hours, the job schedule interval show command will display it with an interval of 1 day 12 hours.

Parameters

[-cluster <Cluster name>] - Cluster

This parameter specifies the name of the cluster on which you want to create an interval schedule. By default, the schedule is created on the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

[-vserver <vserver name>] - Vserver

This parameter specifies the name of the Vserver on which you want to create an interval schedule.

-name <text> - Name

Use this parameter to specify the name of the interval schedule you want to create.

[-days <integer>] - Days

Use this parameter to specify the "days" portion of the schedule's interval. A day is one calendar day.

[-hours <integer>] - Hours

Use this parameter to specify the "hours" portion of the schedule's interval.

[-minutes <integer>] - Minutes

Use this parameter to specify the "minutes" portion of the schedule's interval.

[-seconds <integer>] - Seconds

Use this parameter to specify the "seconds" portion of the schedule's interval.

Examples

The following example creates an interval schedule named rollingdaily that runs six hours after the completion of the previous occurrence of the job:

```
cluster1::> job schedule interval create -name rollingdaily -hours 6
```

Related Links

- job schedule interval show

job schedule interval delete

Delete an interval schedule

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

Description

The `job schedule interval delete` command deletes an interval schedule. Use the job schedule interval show command to display all current interval schedules.

You cannot delete interval schedules that are currently being run. Use the job schedule show-jobs command to

display jobs by schedule.

Parameters

[-cluster <Cluster name>] - Cluster

This parameter specifies the name of the cluster on which you want to delete an interval schedule. By default, the schedule is deleted from the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in switchover state.

[-vserver <vserver name>] - Vserver

This parameter specifies the name of the Vserver on which you want to delete an interval schedule.

-name <text> - Name

Use this parameter with the name of an existing interval schedule to specify the interval schedule you want to delete.

Examples

The following example deletes an interval schedule named rollingdaily:

```
cluster1::> job schedule interval delete -name rollingdaily
```

Related Links

- job schedule interval show
- job schedule show-jobs

job schedule interval modify

Modify an interval schedule

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

Description

The `job schedule interval modify` command modifies an interval schedule. An interval schedule runs jobs at a specified interval after the previous job finishes. Use the job schedule interval show command to display all current interval schedules. See the documentation of job schedule interval show for more information on how interval schedules work.

Modifying one parameter of a schedule's interval does not affect the other parameters. For example, if a schedule's interval is 1 day 12 hours, and you modify the "hours" parameter to 16, the schedule's new interval is 1 day 16 hours. To clear a parameter of the schedule's interval, you must explicitly set that parameter to "0" or "-".

Parameters

[-cluster <Cluster name>] - Cluster

Use this parameter to specify the cluster of an existing interval schedule you want to modify. The local cluster is provided as the default value. In a MetroCluster configuration, the partner cluster can be specified

if the local cluster is in switchover state.

[-vserver <vserver name>] - Vserver

Use this parameter to specify the Vserver of an existing interval schedule you want to modify.

-name <text> - Name

Use this parameter with the name of an existing interval schedule to specify the interval schedule you want to modify.

[-days <integer>] - Days

Use this parameter to specify a different "days" value for the schedule's interval.

[-hours <integer>] - Hours

Use this parameter to specify a different "hours" value for the schedule's interval.

[-minutes <integer>] - Minutes

Use this parameter to specify a different "minutes" value for the schedule's interval.

[-seconds <integer>] - Seconds

Use this parameter to specify a different "seconds" value for the schedule's interval.

Examples

The following example sets the schedule named rollingdaily to run every eight hours:

```
cluster1::> job schedule interval modify -name rollingdaily -hours 8
```

Related Links

- job schedule interval show

job schedule interval show

Show interval schedules

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

Description

The `job schedule interval show` command displays information about interval schedules.

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.

[-cluster <Cluster name>] - Cluster

Selects the interval schedules that match this parameter value.

[-vserver <vserver name>] - Vserver

Selects the interval schedules that match this parameter value. These job schedules can only be used by consumers in that Vserver.

[-name <text>] - Name

Selects the interval schedules that match this parameter value.

[-days <integer>] - Days

Selects the interval schedules that match the day value or range of values you specify.

[-hours <integer>] - Hours

Selects the interval schedules that match the hour value or range of values you specify.

[-minutes <integer>] - Minutes

Selects the interval schedules that match the minute value or range of values you specify.

[-seconds <integer>] - Seconds

Selects the interval schedules that match the second value or range of values you specify.

[-description <text>] - Description

Selects the interval schedules that match the description you specify.

Examples

The following example displays information about all interval schedules:

```
cluster1::> job schedule interval show
Cluster       Vserver     Name        Description
------------- ----------- -----------
-------------------------------------------------------
cluster1      data_vs_1
                          rollingdaily
                                      Every 8h
```

lun commands

lun convert-from-namespace

Transition an existing NVMe namespace into a LUN.

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

Description

This command transitions an existing NVMe namespace in-place into a LUN within a volume in a Vserver. When performing an in-place transition, the application data remains unchanged and only the metadata gets modified.

When using this command, it is necessary that any existing maps to an NVMe subsystem for this namespace be removed first. If any maps exist for the specified namespace, you will receive an error message.

Only NVMe namespaces with a block size of 512 bytes can be transitioned into a LUN. You will receive an error message if the NVMe namespace has a block size different from 512 bytes.

Parameters

`-vserver <Vserver Name>` - Vserver Name

The name of the Vserver containing the NVMe namespace. If only one data Vserver exists, you do not need to specify this parameter.

`-namespace-path <path>` - Path of the NVMe namespace

Specifies the path of the NVMe namespace you want to transition into LUN. Examples of correct namespace paths: *`/vol/vol1/ns1`* and *`/vol/vol1/qtree1/ns1`*

Examples

```
cluster1::> lun convert-from-namespace -vserver vs1 -namespace-path
/vol/vol1/ns1
```

Transitions namespace *`ns1`* in-place to LUN within volume *`vol1`* in Vserver *`vs1`* .

lun create

Create a new LUN

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

Description

This command creates a new LUN of a specific size. You cannot create a LUN at a path that already exists.

You must create LUNs at the root of a volume or qtree. You can not create LUNs in the Vserver root volume.

You might find it useful to provide a meaningful path name for the LUN and containing volume. For example, you might choose a name that describes how the LUN is used, such as the name of the application, the type of data that it stores, or the user accessing the data. Examples are /vol/database/lun0, /vol/finance/lun1, and /vol/bill/lun2.

It is recommended that you distribute LUNs across the cluster.

When you can create a LUN, the size of the LUN could be larger than what you specified. The system generates a message if the size of the LUN is different from what you specified.

By default, when you create a LUN, it is online and it is space-reserved. Use the `lun offline` command to take a LUN offline. When you set space reserved to false, the LUN is non-space reserved.

> For non-space reserved LUNs, write operations to that LUN might fail due to insufficient disk space. As a result, the host application or operating system might crash.

> When you create a LUN from a file, that file cannot be deleted without deleting the LUN itself.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

{ `-path <path>` - LUN Path

Specifies the path of the new LUN. The LUN path cannot contain any files. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name

Specifies the volume that contains the new LUN.

`[-qtree <qtree name>]` - Qtree Name

Specifies the qtree that contains the new LUN.

`-lun <text>` - LUN Name }

Specifies the new LUN name. A LUN name is a case-sensitive name and has the following requirements:

- Must contain one to 255 characters. Spaces are not allowed.
- Can contain the letters A-Z, a-z, numbers 0-9, "-", "_", "}", "{", and ".".

{ `-s, -size <size>` - LUN Size

Specifies the size of the LUN in bytes. You can specify a one-character multiplier suffix:

- c (1 byte)
- w (2 bytes)
- B (512 bytes)
- k (1024 bytes)
- M (k*k bytes)
- G (k*m bytes)

- T (m*m bytes)

[-use-exact-size <true>] - Use Exact Size (privilege: advanced)

Create the LUN using the exact value specified by the `-size` parameter instead of rounding the size to best fit the LUN geometry. Size of the LUN must be a multiple of 512 bytes.

| [-f, -file-path <text>] - File Path }

Creates a LUN using the file path as the source.

| [-foreign-disk <text>] - Foreign Disk Serial number (privilege: advanced) }

LUN is created with the same attributes (size, alignment, bytes per sector and so on) as the specified foreign disk.

[-P, -prefix-size <size>] - Prefix Size (privilege: advanced)

Specifies the size of the prefix stream for the new LUN.

-t, -ostype <LUN Operating System Format> - OS Type

Specifies the OS type for the new LUN. + On an All SAN Array, the following OS types are supported:

- aix - the LUN will be used with AIX.
- hpux - the LUN will be used with HP-UX.
- hyper_v - the LUN will be used with Microsoft Hyper-V.
- linux - the LUN will be used with Linux.
- solaris - the LUN will be used with Solaris slice partitioning.
- vmware - the LUN will be used with VMware VMFS.
- windows - the LUN will be used with Microsoft Windows.
- Xen - the LUN will be used with Xen.

```
+ On all other clusters, the following OS types are supported:
```

- aix - the LUN will be used with AIX.
- hpux - the LUN will be used with HP-UX.
- hyper_v - the LUN will be used with Microsoft Hyper-V.
- linux - the LUN will be used with Linux.
- netware - the LUN will be used with NetWare.
- openvms - the LUN will be used with Open-VMS.
- solaris - the LUN will be used with Solaris slice partitioning.
- solaris_efi - the LUN will be used with Solaris EFI partitioning.
- vmware - the LUN will be used with VMware VMFS.
- windows - the LUN will be used with a Master Boot Record (MBR) partition table on Microsoft Windows 2003 or earlier.
- windows_2008 - the LUN will be used with Microsoft Windows 2008 or later.

- windows_gpt - the LUN will be used with a GUID Partition Type (GPT) partition table on Microsoft Windows.
- xen - the LUN will be used with Xen

[-space-reserve {enabled|disabled}] - Space Reservation

Specifies whether the space reservation setting is *enabled* or *disabled* for the new LUN. If you set the parameter to *enabled* , the LUN is space-reserved. If you set the parameter to *disabled* , the LUN is non-space reserved. The default value is *enabled* .

[-comment <text>] - Comment

A description for the LUN you want to create. If the comment string contains white space, you must enclose the comment string in double quotes. The limit is 254 characters.

[-space-allocation {enabled|disabled}] - Space Allocation

Specifies the value for the space allocation attribute of the LUN. The space allocation attribute determines if the LUN supports the SCSI Thin Provisioning features defined in the Logical Block Provisioning section of the SCSI SBC-3 standard. Specifying enabled for this parameter enables support for the SCSI Thin Provisioning features. Specifying disabled for this parameter disables support for the SCSI Thin Provisioning features. Hosts and file systems that do not support SCSI Thin Provisioning should not enable space allocation. The default is enabled.

[-class {regular|protocol-endpoint|vvol}] - Class

Specifies the class of the new LUN. The class types are:

- regular - the LUN is for normal blocks protocol access. This is the default value.
- protocol-endpoint - the LUN is a vvol protocol endpoint.
- vvol - the LUN is a vvol data LUN.

{ [-qos-policy-group <text>] - QoS Policy Group

This optionally specifies which QoS policy group to apply to the LUN. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to a LUN, the system will not monitor and control the traffic to it.

> If you specify this parameter for a LUN that you want to create from a file and that file belongs to a QoS policy group, ONTAP adds the LUN to the specified policy group and removes the file from its policy group. Both the file and the LUN that you created from the file cannot belong to QoS policy groups.

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

This optionally specifies which QoS adaptive policy group to apply to the LUN. This policy group defines measurable service level objectives (SLOs) and service level agreements (SLAs) that adjust based on the LUN's allocated space or used space.

[-caching-policy <text>] - Caching Policy Name

This optionally specifies the caching policy to apply to the LUN. A caching policy defines how the system caches this volume's data in Flash Cache modules. If a caching policy is not assigned to this LUN, the system uses the caching policy that is assigned to the containing volume or Vserver. If a caching policy is not assigned to the containing volume or Vserver, the system uses the default cluster-wide policy. The available caching policies are:

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

Default caching-policy is auto.

`[-foreground {true|false}]` - Foreground

This parameter specifies whether the operation runs in the foreground. The default setting is true (the operation runs in the foreground). When set to true, the command will not return until the operation completes.

Examples

```
cluster1::> lun create -vserver vs1 -path /vol/vol1/lun1 -size 100M
-ostype linux
```

Creates a 100MB LUN at path *`/vol/vol1/lun1`* in Vserver *`vs1`* . The OS type is Linux, and the state is online.

lun delete

Delete the LUN

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

Description

This command deletes a LUN from a specified Vserver and volume. If the LUN is mapped and online, the force option is required to delete it.

If a LUN is mapped to an initiator group, you can unmap it by using the lun unmap command. If a LUN is online, you take it offline by using the `lun offline` command.

> If you create a LUN from a file, you cannot remove the file while the LUN is linked to it. If you want to remove the file, you must first delete the LUN.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

{ `-path <path>` - LUN Path

Specifies the path of the LUN you want to delete. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name

Specifies the volume that contains the LUN you want to delete.

`[-qtree <qtree name>]` - Qtree Name

Specifies the qtree that contains the LUN you want to delete.

`-lun <text>` - LUN Name }

Specifies the LUN that you want to delete.

`[-f, -force <true>]` - Force Deletion of an Online and Mapped LUN

Force deletion of an online LUN that is mapped to an initiator group.

`[-force-fenced <true>]` - Force Deletion of a Fenced LUN

Force deletion of a LUN that is currently fenced.

`[-foreground {true|false}]` - Foreground

This parameter specifies whether the operation runs in the foreground. The default setting is true (the operation runs in the foreground). When set to true, the command will not return until the operation completes.

`[-disable-smas-proxy <true>]` - Disable SMAS Proxy

Force the modification to run locally without sending operations to the source of the SnapMirror Synchronous relationship. The operation may still fail if not supported on the secondary side in the current state.

Examples

```
cluster1::> lun delete -vserver vs1 -path /vol/vol1/lun1
```

Deletes the LUN at path /vol/vol1/lun1 on Vserver vs1.

Related Links

- lun unmap

lun maxsize

Display the maximum possible size of a LUN on a given volume or qtree.

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

Description

This command returns the maximum size of LUNs for different OS types in a volume or qtree. The command also includes possible maximum size for LUNs with or without snapshots. You can specify the path of the volume or qtree to determine the maximum size of a LUN that you want to create within that volume or qtree.

If you do not specify a path, the command returns the maximum LUN size for each OS type for all volumes and qtrees in a cluster.

The available space in a volume can change over time which means that the size reported by `lun maxsize` can change as well. In addition, the maximum LUN size allowed in a lun resize command may be less than the size reported by `lun maxsize` .

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.

[-vserver <Vserver Name>] - Vserver Name

Specifies the Vserver.

{ [-path <qtree path>] - Volume or Qtree Path

Specifies the path of the root volume or qtree. Examples of correct volume and qtree paths are *`/vol/vol1`* and *`/vol/vol1/qtree1`* .

| [-volume <volume name>] - Volume Name

Specifies the volume that contains the LUN you want to get the maximum size for.

[-qtree <qtree name>] - Qtree Name }

Specifies the qtree that contains the LUN you want to get the maximum size for.

[-t, -ostype <text>] - OS Type

Specifies the OS type for the new LUN. + On an All SAN Array, the following OS types are supported:

- aix - the LUN will be used with AIX.
- hpux - the LUN will be used with HP-UX.
- hyper_v - the LUN will be used with Microsoft Hyper-V.
- linux - the LUN will be used with Linux.
- solaris - the LUN will be used with Solaris slice partitioning.
- vmware - the LUN will be used with VMware VMFS.
- windows - the LUN will be used with Microsoft Windows.
- Xen - the LUN will be used with Xen.

```
+ On all other clusters, the following OS types are supported:
```

- aix - the LUN will be used with AIX.
- hpux - the LUN will be used with HP-UX.
- hyper_v - the LUN will be used with Microsoft Hyper-V.
- linux - the LUN will be used with Linux.
- netware - the LUN will be used with NetWare.
- openvms - the LUN will be used with Open-VMS.
- solaris - the LUN will be used with Solaris slice partitioning.
- solaris_efi - the LUN will be used with Solaris EFI partitioning.
- vmware - the LUN will be used with VMware VMFS.
- windows - the LUN will be used with a Master Boot Record (MBR) partition table on Microsoft Windows 2003 or earlier.
- windows_2008 - the LUN will be used with Microsoft Windows 2008 or later.
- windows_gpt - the LUN will be used with a GUID Partition Type (GPT) partition table on Microsoft Windows.
- xen - the LUN will be used with Xen

[-complete-ss-reserve <size>] - With Complete Snapshot Reserve

Shows the maximum size possible of a LUN if you have the complete snapshot reserve enabled.

[-ss-reserve <size>] - With Snapshot Reserve

Shows the maximum size possible of a LUN if you have snapshot reserve enabled.

[-without-ss-reserve <size>] - Without Snapshot Reserve

Shows the maximum size possible of a LUN if you have no snapshot reserve enabled.

Examples

```
cluster1::> lun maxsize -vserver vs1 -volume vol1 -ostype linux
Virtual                                          Without    With SS
Complete
Server     Volume       Qtree        OS Type  SS Reserve Reserve   SS
Reserve
---------- ------------ ------------ -------- ---------- ---------
----------
vs1        vol1         ""           linux          45MB      45MB
45MB
```

Displays the maximum size of a LUN for the OS type linux.

```
cluster1::> lun maxsize -vserver vs1 -volume vol1
                                                  Without   With SS
Complete
Vserver    Volume       Qtree        OS Type  SS Reserve   Reserve SS
Reserve
---------- ------------ ------------ -------- ---------- ---------
----------
vs1        vol1         ""           aix           178MB     178MB
178MB
             hpux          178MB     178MB      178MB
             hyper_v     172.6MB   172.6MB    172.6MB
             linux         178MB     178MB      178MB
             netware       178MB     178MB      178MB
             openvms       178MB     178MB      178MB
             solaris       178MB     178MB      178MB
             solaris_efi   178MB     178MB      178MB
             windows     172.6MB   172.6MB    172.6MB
             windows_2008
                         172.6MB   172.6MB    172.6MB
             windows_gpt 172.6MB   172.6MB    172.6MB
             xen           178MB     178MB      178MB
12 entries were displayed.
```

Displays the maximum size of LUNs for all OS types on volume vol1.

Related Links

- lun resize

lun modify

Modify a LUN

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

Description

This command modifies LUN attributes. Because LUN modifications can result in data corruption or other problems, we recommend that you call technical support if you are unsure of the possible consequences of modifying a LUN.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

{ -path <path> - LUN Path

Specifies the path for the LUN you want to modify. Examples of correct LUN paths: */vol/vol1/lun1* and */vol/vol1/qtree1/lun1*

| -volume <volume name> - Volume Name

Specifies the volume for the LUN you want to modify.

-qtree <qtree name> - Qtree Name

Specifies the qtree for the LUN you want to modify.

-lun <text> - LUN Name }

Specifies the name for the LUN you want to modify. A LUN name is a case-sensitive name and has the following requirements:

- Must contain one to 255 characters. Spaces are not allowed.
- Can contain the letters A through Z, a through z, numbers 0 through 9, hyphen (-), underscore (_), right bracket (}), left bracket ({) and period (.).
- Must start with a letter or number.

[-space-reserve {enabled|disabled}] - Space Reservation

Specifies whether the space reservation setting is enabled or disabled for a LUN. If you set the parameter to *enabled* , the LUN is space-reserved. If you set the parameter to *disabled* , the LUN is non-space reserved. The default is *enabled* .

{ [-serial <text>] - Serial Number

Specifies the serial number for the LUN you want to modify.

The LUN serial number is a twelve-character alphanumeric string containing one or more of the following:

- upper- and lower-case letters
- numbers
- the characters: , <, >, /, -, #, $, %, *, +, =, ?, @, [, !,], ^, ~

Some of the characters that are valid in a LUN serial number also have special meaning to the cluster shell command line:

- The question mark (?) activates the command line active help. In order to type a question mark as part of a LUN's serial number, it is necessary to disable active help with the command set -active-help false . Active help can later be re-enabled with the command set -active-help true .
- The number sign (#) indicates the beginning of a comment to the command line and will cause the remainder of the line to be ignored. To avoid this, enclose the serial number in double quotes (").

Alternatively, the `-serial-hex` parameter can be used to set the LUN serial number specifying the serial number encoded in hexadecimal form.

| [-serial-hex <Hex String>] - Serial Number (Hex) }

Specifies the serial number, encoded in hexadecimal form, for the LUN you want to modify. See the description of the `-serial` parameter for additional details.

[-comment <text>] - Comment

Specifies the comment for the LUN you want to modify.

[-space-allocation {enabled|disabled}] - Space Allocation

Specifies the new value for the space allocation attribute of the LUN. The space allocation attribute determines if the LUN supports the SCSI Thin Provisioning features defined in the Logical Block Provisioning section of the SCSI SBC-3 standard.

Specifying *enabled* for this parameter enables support for the SCSI Thin Provisioning features.

Specifying *disabled* for this parameter disables support for the SCSI Thin Provisioning features.

Hosts and file systems that do not support SCSI Thin Provisioning should not enable space allocation.

[-state {online|offline|nvfail|space-error|foreign-lun-error}] - State

Specifies the administrative state of a LUN. The options are:

- online
- offline

{ [-device-legacy-id <integer>] - Device Legacy ID

Specifies the device legacy ID for the LUN you want to modify.

| [-device-binary-id <text>] - Device Binary ID

Specifies the device binary ID for the LUN you want to modify.

| [-clear-binary-id <true>] - Clear Device Binary ID }

Clears the binary format of the optional device ID.

{ [-device-text-id <text>] - Device Text ID

Specifies the device text ID for the LUN you want to modify.

| [-clear-text-id <true>] - Clear Device Text ID }

Clears the text format of the optional device ID.

{ [-qos-policy-group <text>] - QoS Policy Group

This optionally specifies which QoS policy group to apply to the lun. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to a lun, the system will not monitor and control the traffic to it. To remove this lun from a policy group, enter the reserved keyword "none".

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

This optional parameter specifies which QoS adaptive policy group to apply to the LUN. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the LUN’s allocated space or used space. To remove this LUN from an adaptive policy group, enter the reserved keyword "none".

[-caching-policy <text>] - Caching Policy Name

This optionally specifies the caching policy to apply to the LUN. A caching policy defines how the system caches this volume’s data in Flash Cache modules. If a caching policy is not assigned to this LUN, the

system uses the caching policy that is assigned to the containing volume or Vserver. If a caching policy is not assigned to the containing volume or Vserver, the system uses the default cluster-wide policy. The available caching policies are:

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

Default caching-policy is auto.

`[-disable-smas-proxy <true>]` - Disable SMAS Proxy (privilege: advanced)

Force the modification to run locally without sending the operation to the source of the SnapMirror Synchronous relationship. The operation may still fail if not supported on the secondary side in the current state.

Examples

```
cluster1::> lun modify -path /vol/vol1/lun1 -space-reserve disable
```

Disables the space reserve attribute for LUN /vol/vol1/lun1.

```
cluster1::> lun modify -path /vol/vol1/lun1 -state offline
```

Takes the LUN /vol/vol1/lun1 offline.

```
cluster1::> lun modify -path /vol/vol1/lun1 -comment "new comment"
```

Adds the comment "new comment" to the LUN /vol/vol1/lun1.

Related Links

- set

lun move-in-volume

Move a LUN within a volume

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

Description

This command moves a LUN to a new path in the same volume or renames a LUN. If you are organizing LUNs in a qtree, the command moves a LUN from one qtree to another. You can perform a LUN move while the LUN is online and serving data. The process is non-disruptive. Use the lun move start command to move a LUN to a different volume within the same Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

{ `-path <path>` - LUN Path

Specifies the path of the LUN you want to move. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

| `-volume <volume name>` - Volume Name

Specifies the volume of the LUN you want to move.

`[-qtree <qtree name>]` - Qtree Name

Specifies the qtree of the LUN you want to move.

`-lun <text>` - LUN Name }

Specifies the name of the LUN that you want to move.

{ `-new-path <path>` - New LUN Path

Specifies the new path of the LUN. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

| `[-new-qtree <qtree name>]` - New Qtree Name

Specifies the new qtree name that you want to move the LUN to.

`-new-lun <text>` - New LUN Name }

Specifies the new name of the LUN.

Examples

```
cluster1::> lun move-in-volume -vserver vs1 -volume vol1 -lun lun1 -new
-lun newlun1
```

Renames *`lun1`* to *`newlun1`* on Vserver *`vs1`* and volume *`vol1`* .

```
cluster1::> lun show -vserver vs1 -volume vol1
server   Path                            State   Mapped   Type
Size
-------- ------------------------------- ------- -------- --------
--------
s1       /vol/vol1/A/lun1                online  mapped   linux
10MB

luster1::> lun move-in-volume -vserver vs1 -path /vol/vol1/A/lun1 -new
-path /vol/vol1/B/lun1

luster1::> lun show -vserver vs1 -volume vol1
server   Path                            State   Mapped   Type
Size
-------- ------------------------------- ------- -------- --------
--------
s1       /vol/vol1/B/lun1                online  mapped   linux
10MB
```

Moves LUN *lun1* from qtree *A* to qtree *B* on volume *vol1* .

Related Links

- lun move start

lun resize

Changes the size of the LUN to the input value size.

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

Description

This command resizes an existing LUN.

If the value specified by the `-size` parameter is larger than the current size of the LUN, the LUN size will be increased to at least the requested size by appending new empty blocks to the end of the LUN. After resizing, follow the procedures provided by your host operating system to discover the new size of the LUN and expand the filesystem on the LUN.

If the value specified by the `-size` parameter is smaller than the current size of the LUN, the LUN size will be reduced to no less than the requested size by truncating blocks at the end of the LUN beyond the new size. All data in the truncated blocks will be lost. Prior to reducing the size of a LUN, you must follow the procedure provided by the host operating system to resize and migrate the filesystem on the LUN off the blocks that will be truncated. Not all operating systems and filesystems support shrinking an existing filesystem.

You will receive an error message warning of the potential for data loss if the command would reduce the size of the LUN. Once the host filesystem has been migrated off the blocks to be truncated, use the `-force`

parameter to allow the resize to proceed.

> Consider taking a snapshot of the volume before reducing the size of the LUN. After the LUN has been resized and the contents verified, you can delete the snapshot.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

{ `-path <path>` - LUN Path

Specifies the path of the LUN that you want to resize. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name

Specifies the volume that contains the LUN that you want to resize.

`[-qtree <qtree name>]` - Qtree Name

Specifies the qtree that contains the LUN that you want to resize.

`-lun <text>` - LUN Name }

Specifies the LUN name that you want to resize.

`[-f, -force <true>]` - Force Reduce LUN Size

Specifies the command may reduce the size of the specified LUN. If the specified `-size` is smaller than the existing size of the LUN and this parameter is not specified, the command will fail with an error instead of potentially destroying data. See the overall command description above for more information concerning reducing the size of a LUN.

`[-size <size>]` - New Size

Specifies the requested new size of the LUN. The actual size after resizing may be slightly larger than requested. If the value is prefixed with `+` or `-` , the new size will be calculated as the existing size plus or minus the value provided. The following multipliers are recognized as suffixes:

- B - Value specifies number of blocks (512 bytes)
- k - Value specifies number of kilobytes (1024 bytes)
- M - Value specifies number of megabytes (1024 kilobytes)
- G - Value specifies number of gigabytes (1024 megabytes)
- T - Value specifies number of terabytes (1024 gigabytes)

Examples

```
cluster1::> lun resize -vserver vs1 -path /vol/vol1/lun1 -size 500M
```

Resizes the LUN /vol/vol1/lun1 in Vserver vs1 to 500 MB. If the LUN is already larger than 500 MB, this command will fail.

```
cluster1::> lun resize -vserver vs1 -path /vol/vol1/lun1 -size +50M
```

Adds 50 MB of space to LUN /vol/vol1/lun1

```
cluster1::> lun resize -vserver vs1 -path /vol/vol1/lun1 -size -10m

Error: command failed: Reducing LUN size without coordination with the
host system
may cause permanent data loss or corruption. Use the force flag to allow
LUN size reduction.

cluster1::> lun resize -path /vol/vol1/lun1 -size -100M -force
```

Reduces the LUN /vol/vol1/lun1 by 100 MB.

lun show

Display a list of LUNs

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

Description

The command displays information for LUNs. Use the `instance` parameter to display additional LUN details, such as serial number and space-reservation settings.

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.

[-vserver <Vserver Name>] - Vserver Name

Selects the LUNs that match this parameter value.

{ [-path <path>] - LUN Path

Selects the LUNs that match this parameter value. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| [-volume <volume name>] - Volume Name

Selects the LUNs that match this parameter value.

[-qtree <qtree name>] - Qtree Name

Selects the LUNs that match this parameter value.

[-lun <text>] - LUN Name }

Selects the LUNs that match this parameter value.

[-s, -size <size>] - LUN Size

Selects the LUNs that match this parameter value.

[-P, -prefix-size <size>] - Prefix Size (privilege: advanced)

Selects the LUNs that match the prefix stream size that you specify.

[-t, -ostype <LUN Operating System Format>] - OS Type

Selects the LUNs that match this parameter value. The operating system types are:

- *aix* - the LUN stores AIX data.
- *hpux* - the LUN stores HP-UX data.
- *hyper_v* - the LUN stores Windows Server 2008 or Windows Server 2012 Hyper-V data
- *linux* - the LUN stores a Linux raw disk without a partition table.
- *netware* - the LUN stores NetWare data.
- *openvms* - the LUN stores Open-VMS data
- *solaris* - the LUN stores Solaris raw disk in a single-slice partition.
- *solaris_efi* - the LUN stores Solaris_EFI data.
- *vmware* - the LUN stores VMware data
- *windows* - the LUN stores a raw disk type in a single-partition Windows disk using the Master Boot Record (MBR) partitioning style.
- *windows_gpt* - the LUN stores Windows data using the GUID Partition Type (GPT) partitioning style.
- *windows_2008* - the LUN stores Windows data for Windows 2008 or later systems.
- *xen* - the LUN stores Xen data

[-space-reserve {enabled|disabled}] - Space Reservation

Selects the LUNs that match this parameter value. A value of *enabled* selects LUN that are space-reserved. A value of *disabled* selects LUNs that are non-space reserved.

[-serial <text>] - Serial Number

Selects the LUNs that match this parameter value.

The LUN serial number is a twelve-character alphanumeric string containing one or more of the following:

- upper- and lower-case letters
- numbers
- the characters: , <, >, /, -, #, $, %, *, +, =, ?, @, [, !,], ^, ~

Some of the characters that are valid in a LUN serial number also have special meaning to the cluster shell

command:

- The question mark (?) activates the command line active help. In order to type a question mark as part of a LUN's serial number, it is necessary to disable active help with the command set -active-help false . Active help can later be re-enabled with the command set -active-help true .
- The number sign (#) indicates the beginning of a comment to the command line and will cause the remainder of the line to be ignored. To avoid this, enclose the serial number in double quotes (").
- The less than (<), greater than (>), asterisk (*), and exclamation point (!) influence the query behavior of the command. To use these as characters in a LUN's serial query, you must first press escape (ESC). To use these characters to influence the query, enclose the serial number, or partial serial number, in double quotes (") and apply <, >, *, or !, outside of the double quotes.

Alternatively, the `-serial-hex` parameter can be used to select LUNs using the serial number encoded in hexadecimal form.

`[-serial-hex <Hex String>]` - Serial Number (Hex)

Selects the LUNs that match this parameter value. This parameter applies to the LUN serial number encoded in hexadecimal form. See the description of the `-serial` parameter for additional details.

`[-comment <text>]` - Comment

Selects the LUNs that match this parameter value.

`[-space-reserve-honored {true|false}]` - Space Reservations Honored

Selects the LUNs that match this parameter value. A value of *`true`* selects LUNs that have their space reservation honored by the container volume. A value of *`false`* displays the LUNs that are non-space reserved.

`[-space-allocation {enabled|disabled}]` - Space Allocation

Selects the LUNs that match this parameter value. The space allocation attribute determines if the LUN supports the SCSI Thin Provisioning features defined in the Logical Block Provisioning section of the SCSI SBC-3 standard.

Specifying *`enabled`* for this parameter selects LUNs with support enabled for the SCSI Thin Provisioning features.

Specifying *`disabled`* for this parameter selects LUNs with support disabled for the SCSI Thin Provisioning features.

Hosts and file systems that do not support SCSI Thin Provisioning should not enable space allocation.

`[-container-state {online|aggregate-offline|volume-offline|error}]` - LUN Container State (privilege: advanced)

Selects the LUNs that match this parameter value. The container states are:

- *`online`* - The LUN's aggregate and volume are online.
- *`aggregate-offline`* - The LUN's aggregate is offline.
- *`volume-offline`* - The LUN's volume is offline.
- *`error`* - An error occurred accessing the LUN's volume.

[-state {online|offline|nvfail|space-error|foreign-lun-error}] - State

Selects the LUNs that match this parameter value. The states are:

- *online* - The LUN is online.
- *offline* - The LUN is administratively offline, or a more detailed offline reason is not available.
- *foreign-lun-error* - The LUN has been automatically taken offline due to a media error on the associated foreign LUN.
- *nvfail* - The LUN has been automatically taken offline due to an NVRAM failure.
- *space-error* - The LUN has been automatically taken offline due to insufficient space.

[-uuid <UUID>] - LUN UUID

Selects the LUNs that match this parameter value.

[-mapped {mapped|unmapped}] - Mapped

Selects the LUNs that match this parameter value. A value of `mapped` selects the LUNs that are mapped to an initiator group.

[-block-size {512|4KB}] - Physical Size of Logical Block

Selects the LUNs that match this parameter value.

[-device-legacy-id <integer>] - Device Legacy ID

Selects the LUNs that match this parameter value.

[-device-binary-id <text>] - Device Binary ID

Selects the LUNs that match this parameter value.

[-device-text-id <text>] - Device Text ID

Selects the LUNs that match this parameter value.

[-read-only {true|false}] - Read Only

Selects the LUNs that match this parameter value.

[-restore-inaccessible {true|false}] - Fenced Due to Restore

Selects the LUNs that match the state you specify. A value of *true* means that a LUN is fenced for I/O and management due to a restore operation.

[-size-used <size>] - Used Size

Selects the LUNs that match this parameter value.

[-max-resize-size <size>] - Maximum Resize Size

Selects the LUNs that match this parameter value.

[-creation-timestamp <MM/DD/YYYY HH:MM:SS>] - Creation Time

Selects the LUNs that match this parameter value.

[-class {regular|protocol-endpoint|vvol}] - Class

Selects the LUNs that match this parameter value.

[-node <nodename>] - Node Hosting the LUN

Selects the LUNs that match this parameter value.

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

Selects the LUNs that match this parameter value.

A policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to a LUN, the system will not monitor and control the traffic to it.

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

Selects the LUNs that match this parameter value.

An adaptive policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the adaptive policy group is associated.

[-caching-policy <text>] - Caching Policy Name

Display the LUNs that match the specified cache.

A caching policy defines the caching behavior of this LUN at the Flash Cache level. If a caching policy is not assigned to this LUN, the system uses the caching policy that is assigned to the containing volume or Vserver. If a caching policy is not assigned to the containing volume or Vserver, the system uses the default cluster-wide policy. The available caching policies are:

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

Default caching-policy is auto.

[-is-clone {true|false}] - Clone

Selects the LUNs that match this parameter value.

[-is-clone-autodelete-enabled {true|false}] - Clone Autodelete Enabled

Selects the LUNs that match this parameter value.

[-inconsistent-import {true|false}] - Inconsistent Import

Selects the LUNs that match this parameter value. A value of *true* means that the import of this LUN is incomplete.

[-serial-7-mode <text>] - 7-mode Serial Number (privilege: advanced)

Selects the LUNs that match this parameter value.

LUNs transitioned from ONTAP 7-Mode are assigned new serial numbers for use with ONTAP. The original 7-Mode serial number is displayed in this field for reference.

[-application <text>] - Application

Selects the LUNs that are part of an application that matches the parameter value.

[-include-offline-containers <true>] - Include LUNs on Offline Volumes and Aggregates (privilege: advanced)

If true, include available information for LUNs in offline aggregates and offline volumes in the output. By default, LUNs in offline aggregates and offline volumes are excluded from the output.

[-zrto-vol-consensus {true|false}] - ZRTO Volume consensus

Selects the LUNs that match the specified value. The status indicates whether the volume containing the LUN in a SnapMirror Synchronous relationship can process operations.

[-consistency-group <text>] - Consistency Group Name

Selects the LUNs that are associated with the specified consistency group.

Examples

The following example displays details of the LUN at path /vol/vol1/lun1 in Vserver vs1.

```
cluster1::> lun show -vserver vs1 -path /vol/vol1/lun1 -instance
              Vserver Name: vs1
                  LUN Path: /vol/vol1/lun1
               Volume Name: vol1
                Qtree Name: ""
                  LUN Name: lun1
                  LUN Size: 10MB
                   OS Type: linux
         Space Reservation: disabled
             Serial Number: wCVt1]IlvQWv
       Serial Number (Hex): 77435674315d496c76515776
                   Comment: new comment
Space Reservations Honored: false
          Space Allocation: disabled
                     State: offline
                  LUN UUID: 76d2eba4-dd3f-494c-ad63-1995c1574753
                    Mapped: mapped
                Block Size: 512
          Device Legacy ID: -
          Device Binary ID: -
            Device Text ID: -
                 Read Only: false
     Fenced Due to Restore: false
                 Used Size: 5MB
       Maximum Resize Size: 64.00GB
             Creation Time: 9/14/2016 13:55:09
                     Class: regular
      Node Hosting the LUN: node1
          QoS Policy Group: -
       Caching Policy Name: -
                     Clone: false
  Clone Autodelete Enabled: false
       Inconsistent Import: false
               Application: -
```

The following example displays information for the LUN with serial number 1r/wc+9Cpbls:

```
cluster1::> lun show -serial 1r/wc+9Cpbls
Vserver   Path                            State   Mapped   Type
Size
--------- ------------------------------- ------- -------- --------
--------
vs1       /vol/vol2/lun1                  online  mapped   linux
10MB
```

The following example displays all the LUNs on Vserver vs1 and volume vol1:

```
cluster1::> lun show -vserver vs1 -volume vol1
Vserver   Path                            State   Mapped   Type
Size
--------- ------------------------------- ------- -------- --------
--------
vs1       /vol/vol1/lun1                  offline mapped   linux
10MB
vs1       /vol/vol1/lun2                  online  mapped   windows
47.07MB
2 entries were displayed.
```

Related Links

- set

lun bind commands

lun bind create

Bind a VVol LUN to a protocol endpoint

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

Description

This command creates a new binding between a protocol endpoint and a vvol LUN. If a binding between the specified endpoint and vvol already exists, the reference count for the binding is incremented by one.

> For optimal results, the protocol endpoint and vvol must be hosted by the same node in the cluster.

Parameters

`-vserver <Vserver Name>` - Vserver name (privilege: advanced)

Specifies the name of the Vserver.

`-protocol-endpoint-path <path>` - Protocol Endpoint (privilege: advanced)

Specifies the path to the protocol endpoint. The specified LUN must already exist and be of class "protocol-endpoint". Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

`-vvol-path <path>` - VVol Path (privilege: advanced)

Specifies the path to the vvol. The specified LUN must already exist and be of the class "vvol". Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

Examples

```
cluster1::*> lun bind create -vserver vs1 -protocol-endpoint-path
/vol/VV1/PE1 -vvol-path /vol/VV3/234ace
```

Bind the vvol /vol/VV3/234ace to the protocol endpoint /vol/VV1/PE1 in Vserver vs1.

lun bind destroy

Unbind a VVol LUN from a protocol endpoint

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

Description

Decrement the reference count of the binding between a protocol endpoint and vvol LUN. If the resulting reference count is zero, the binding is removed.

Parameters

`-vserver <Vserver Name>` - Vserver name (privilege: advanced)

Specifies the Vserver.

`-protocol-endpoint-path <path>` - Protocol Endpoint (privilege: advanced)

Specifies the path of the protocol endpoint LUN. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

`-vvol-path <path>` - VVol Path (privilege: advanced)

Specifies the path of the vvol LUN. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

`[-force <true>]` - If true, unbind the Vvol completely even if the current reference count is greater than 1. The default is false. (privilege: advanced)

Completely remove the specified binding, regardless of the current reference count.

Examples

```
cluster1::*> lun bind destroy -protocol-endpoint-path /vol/VV2/PE2 -vvol
-path /vol/VV2/30dfab -vserver vs1
```

Remove the binding between the vvol /vol/VV2/30dfab and the protocol endpoint /vol/VV2/PE2 on Vserver vs1.

lun bind show

Show list of Vvol bindings

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

Description

Shows the configured VVol to protocol endpoint bindings.

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.

[-vserver <Vserver Name>] - Vserver Name (privilege: advanced)

Selects the bindings that match this parameter value.

[-protocol-endpoint-msid <integer>] - PE MSID (privilege: advanced)

Selects the bindings that match this parameter value.

[-protocol-endpoint-vdisk-id <text>] - PE Vdisk ID (privilege: advanced)

Selects the bindings that match this parameter value.

[-vvol-msid <integer>] - VVol MSID (privilege: advanced)

Selects the bindings that match this parameter value.

[-vvol-vdisk-id <text>] - VVol Vdisk ID (privilege: advanced)

Selects the bindings that match this parameter value.

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

Selects the bindings that match this parameter value.

[-protocol-endpoint-path <path>] - Protocol Endpoint (privilege: advanced)

Selects the bindings that match this parameter value. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

[-protocol-endpoint-node <nodename>] - PE Node (privilege: advanced)

Selects the bindings that match this parameter value.

[-vvol-path <path>] - VVol (privilege: advanced)

Selects the bindings that match this parameter value. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

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

Selects the bindings that match this parameter value.

[-secondary-lun <Hex 64bit Integer>] - Secondary LUN (privilege: advanced)

Selects the bindings that match this parameter value.

`[-is-optimal {true|false}]` - Optimal binding (privilege: advanced)

Selects the bindings that match this parameter value.

`[-reference-count <integer>]` - Reference Count (privilege: advanced)

Selects the bindings that match this parameter value.

Examples

```
cluster1::*> lun bind show -vserver vs1
Vserver        Protocol Endpoint                  Node
                   Vvol LUN                            Secondary LUN
Optimal?
-------------- ------------------------------------ --------------
--------
vs1            /vol/VV1/PE1                       cluster-node1
                   /vol/VV2/30dfab                       d20000010000 false
                   /vol/VV3/234ace                       d20000020000 true
                   /vol/VV3/234acf                       d20000030000 true
               /vol/VV2/PE2                       cluster-node2
                   /vol/VV2/30dfab                       d20000010000 true
4 entries were displayed.
```

The example above displays all the LUN bindings on Vserver vs1.

lun copy commands

lun copy cancel

Cancel a LUN copy operation before the new LUN has been created

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

Description

The `lun copy cancel` command cancels an ongoing LUN copy operation prior to creation of the new LUN. The command fails if the LUN already exists at the destination path; in that case, use the lun delete command to delete the LUN at the destination path.

All data transfers will be halted.

> This is an advanced command because the preferred way to cancel a LUN copy operation is to wait until the new LUN becomes visible, and then use the lun delete command to delete the LUN.

Parameters

{ -vserver <Vserver Name> - Vserver Name (privilege: advanced)

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path (privilege: advanced)

Specifies the path of the new LUN. Examples of correct LUN paths: */vol/vol1/lun1* and */vol/vol1/qtree1/lun1*

Examples

```
cluster1::*> lun copy cancel -vserver vs1 -destination-path /vol/vol2/lun2
```

Cancels the copy operation identified by Vserver *vs1* and destination path */vol/vol2/lun2* .

Related Links

- lun delete

lun copy modify

Modify an ongoing LUN copy operation

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

Description

The `lun copy modify` command modifies the maximum throughput of an ongoing copy operation.

Parameters

{ -vserver <Vserver Name> - Vserver Name

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path

Specifies the path of the new LUN. Examples of correct LUN paths: */vol/vol1/lun1* and */vol/vol1/qtree1/lun1*

-max-throughput {<integer>[KB|MB|GB|TB|PB]} - Maximum Transfer Rate (per sec)

Specifies the maximum amount of data, in bytes, that can be transferred per second in support of this operation. This mechanism can be used to throttle a transfer, to reduce its impact on the performance of the source and destination nodes.

> The specified value will be rounded up to the nearest megabyte.

Examples

```
cluster1::> lun copy modify -vserver vs1 -destination-path /vol/vol2/lun2
-max-throughput 25MB
```

Modifies the maximum throughput for the ongoing copy job identified by Vserver *vs1* and destination path */vol/vol2/lun2* to 25 MB/sec.

lun copy pause

Pause an ongoing LUN copy operation

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

Description

The `lun copy pause` command pauses an ongoing copy operation. Use the lun copy resume command to resume the copy operation.

Parameters

{ -vserver <Vserver Name> - Vserver Name

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path

Specifies the path of the new LUN. Examples of correct LUN paths: */vol/vol1/lun1* and */vol/vol1/qtree1/lun1*

Examples

```
cluster1::> lun copy pause -vserver vs1 -destination-path /vol/vol2/lun2
```

Pauses the ongoing copy operation identified by Vserver *vs1* and destination path */vol/vol2/lun2* .

Related Links

- lun copy resume

lun copy resume

Resume a paused LUN copy operation

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

Description

The `lun copy resume` command resumes a paused copy operation.

Parameters

{ -vserver <Vserver Name> - Vserver Name

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path

Specifies the path of the new LUN. Examples of correct LUN paths: */vol/vol1/lun1* and */vol/vol1/qtree1/lun1*

Examples

```
cluster1::> lun copy resume -vserver vs1 -destination-path /vol/vol2/lun2
```

Resumes the paused copy operation identified by Vserver *vs1* and destination path */vol/vol2/lun2* .

lun copy show

Display a list of LUNs currently being copied

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

Description

The `lun copy show` command shows information about LUNs currently being copied in the cluster.

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.

[-vserver <Vserver Name>] - Destination Vserver Name

Selects LUN copy operations that match this parameter value.

[-destination-path <path>] - Destination Path

Selects LUN copy operations that match this parameter value.

[-source-vserver <vserver name>] - Source Vserver Name

Selects LUN copy operations that match this parameter value.

[-source-path <path>] - Source Path

Selects LUN copy operations that match this parameter value.

[-source-snapshot <snapshot name>] - Source Snapshot Name

Selects LUN copy operations that match this parameter value.

[-is-promoted-early {true|false}] - Is Destination Promoted Early

Selects LUN copy operations that match this parameter value.

[-max-throughput {<integer>[KB|MB|GB|TB|PB]}] - Maximum Transfer Rate (per sec)

Selects LUN copy operations that match this parameter value.

[-job-status {Preparing|Allocation-Map|Data|Destroying|Paused-Manual|Paused-Error|Complete|Destroyed|Terminated-Manual}] - LUN Copy Status

Selects LUN copy operations that match this parameter value. The possible values are:

- `Preparing` - the LUN copy job is in Preparing status.
- `Allocation-Map` - the LUN copy job is in Allocating status.
- `Data` - the LUN copy job is in Moving Data status.
- `Destroying` - the LUN copy job is in Destroying status.
- `Paused-Manual` - the LUN copy job is in Manually Paused status.
- `Paused-Error` - the LUN copy job is in Paused By Error status.
- `Complete` - the LUN copy job is in Complete status.
- `Destroyed` - the LUN copy job is in Destroyed status.
- `Terminated-Manual` - the LUN copy job is in Manually Terminated status.

[-progress-percent <percent>] - LUN Copy Progress (%)

Selects LUN copy operations that match this parameter value.

[-elapsed-time <time_interval>] - Elapsed Time

Selects LUN copy operations that match this parameter value.

[-cutover-time <time_interval>] - Cutover Time

Selects LUN copy operations that match this parameter value.

[-is-snapshot-fenced {true|false}] - Is Snapshot Fenced

Selects LUN copy operations that match this parameter value.

[-is-destination-ready {true|false}] - Is Destination Ready

Selects LUN copy operations that match this parameter value.

[-last-failure-reason <text>] - Last Failure Reason

Selects LUN copy operations that match this parameter value.

Examples

```
cluster1::> lun copy show
Vserver   Destination Path                Status          Progress
--------- ------------------------------- --------------- --------
vs1       /vol/vol2/lun1                  Data            35%
vs1       /vol/vol2/lun2                  Complete        100%
2 entries were displayed.
```

The example above displays information about all the LUN copy operations in the cluster.

```
cluster1::> lun copy show -vserver vs1 -destination-path /vol/vol2/lun1
-instance
Destination Vserver Name: vs1
               Destination Path: /vol/vol2/lun1
            Source Vserver Name: vs1
                    Source Path: /vol/vol1/lun1
           Source Snapshot Name: -
  Is Destination Promoted Early: false
Maximum Transfer Rate (per sec): 0B
                LUN Copy Status: Data
          LUN Copy Progress (%): 35%
                   Elapsed Time: 145s
            Cutover Time (secs): 0s
             Is Snapshot Fenced: true
           Is Destination Ready: true
            Last Failure Reason: -
```

The example above displays all information about the LUN being copied to `/vol/vol2/lun1` in Vserver `vs1` .

lun copy start

Start copying a LUN from one volume to another within a cluster

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

Description

The `lun copy start` command initiates copying of a LUN from one volume to another. The destination volume can be located in the same Vserver as the source volume (intra-Vserver copy) or in a different Vserver (inter-Vserver).

> A cluster administrator must first create a Vserver peering relationship using vserver peer create before initiating an inter-Vserver LUN copy operation.

Parameters

`-vserver <Vserver Name>` - Destination Vserver Name

Specifies the name of the Vserver that will host the new LUN.

| `-destination-path <path>` - Destination Path

Specifies the path of the new LUN. Examples of correct LUN paths: `/vol/vol1/lun1` and `/vol/vol1/qtree1/lun1`

`-source-path <path>` - Source Path }

Specifies the path of the source LUN. Examples of correct LUN paths: `/vol/vol1/lun1` ,

/vol/vol1/qtree1/lun1 , and */vol/vol1/.snapshot/qtree/lun1*

[-source-vserver <vserver name>] - Source Vserver Name

Optionally specifies the name of the Vserver hosting the LUN to be copied.

If this parameter is not specified, it is assumed that an intra-Vserver copy operation is being initiated. The source volume is expected to be in the same Vserver as the destination volume.

[-promote-early <true>] - Promote Early

Optionally specifies that the destination LUN needs to be promoted early.

If the destination is promoted early, the new LUN will be visible immediately. However, snapshots of the volume containing the new LUN cannot be taken until the LUN copy operation reaches 'Moving Data' status.

If the destination is promoted late, the new LUN will be visible only after it has been fully framed. However, the LUN copy job will not block the creation of snapshots of the volume containing the new LUN.

If this parameter is not specified, the destination LUN will be promoted late.

[-max-throughput {<integer>[KB|MB|GB|TB|PB]}] - Maximum Transfer Rate (per sec)

Optionally specifies the maximum amount of data, in bytes, that can be transferred per second in support of this operation. This mechanism can be used to throttle a transfer, to reduce its impact on the performance of the source and destination nodes.

If this parameter is not specified, throttling is not applied to the data transfer.

> The specified value will be rounded up to the nearest megabyte.

Examples

```
cluster1::> lun copy start -vserver vs2 -destination-path /vol/vol2/lun2
-source-vserver vs1 -source-path /vol/vol1/lun1
```

Starts an inter-Vserver copy of LUN *lun1* from volume *vol1* in Vserver *vs1* to *lun2* on volume *vol2* in Vserver *vs2* .

```
cluster1::> lun copy start -vserver vs1 -destination-path /vol/vol2/lun2
-source-path /vol/vol1/lun1
```

Starts an intra-Vserver copy of LUN *lun1* from volume *vol1* in Vserver *vs1* to *lun2* on volume *vol2* in Vserver *vs1* .

```
cluster1::> lun copy start -vserver vs1 -destination-path /vol/vol2/lun2
-source-path /vol/vol1/.snapshot/snap1/lun1
```

Starts an intra-Vserver copy of LUN *lun1* from snapshot *snap1* of volume *vol1* in Vserver *vs1* to *lun2* on

volume *vol2* in Vserver *vs1* .

Related Links

- vserver peer create

lun igroup commands

lun igroup add

Add initiators to an initiator group

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

Description

This command adds initiators to an existing initiator group (igroup). You can add an initiator to an initiator group only if there are no LUN mapping conflicts. Mapping conflicts occur when an initiator is already paired with a LUN. If you attempt to run this command and there are LUN mapping conflicts, the command returns an error.

An initiator cannot be a member of two igroups of different OS types. For example, if you have an initiator that belongs to a Solaris igroup, the command does not allow you to add this initiator to an AIX igroup.

When you add FCP initiators, you can specify an alias instead of the initiator's World Wide Port Name (WWPN) or the iSCSI Qualified name (IQN).

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the initiator group to which you want to add a new initiator.

`[-initiator <text>,…]` - Initiators

Specifies the initiator that you want to add. You can specify the WWPN, IQN, or alias of the initiator.

`[-child-igroups <text>,…]` - Child Igroups

Specifies the child initiator groups you want to add. The initiator group inherits the initiators of any child igroups. An initiator group can have either initiators or child initiator groups, but not both simultaneously. Initiator groups can be nested to match host cluster hierarchies. If the host cluster then changes, initiators only need to be updated in the child igroups and the parent igroups inherit the changes.

`[-proximal-vserver {local|peer}]` - Proximal Vserver (privilege: advanced)

Specifies a Vserver to add to any initiators being added to an initiator group. *local* refers to the initiator group's Vserver. *peer* refers to the initiator group's replication peer. *proximal-vserver* can only be set for replicated initiator groups.

Examples

```
cluster1::> lun igroup add -vserver vs1 -igroup ig1 -initiator iqn.1992-
08.com.mv.mvinitiator
```

Adds the initiator iqn.1992-08.com.mv.mvinitiator to the initiator group ig1 on Vserver vs1.

lun igroup bind

Bind an existing initiator group to a given portset

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

Description

This command binds an initiator group to a portset so the host knows which LIFs or TPGs to access. When you bind a portset to an igroup, the host knows which iSCSI or FCP LIF to access. If you do not bind an igroup to a portset, and you map a LUN to the igroup, then the initiators in the igroup can access the LUN on any port on the Vserver.

The initiator group cannot be bound to another portset when you use this command. If you attempt to bind a portset to an initiator group that is already bound to an existing portset, the command returns an error. You can only bind an initiator group to one portset at a time.

If the initiator group is bound, use the lun igroup unbind command to unbind the initiator group from the portset. After the initiator group is unbound, you can bind it to another portset.

You can only bind an initiator group to a non-empty portset.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the initiator group that you want to bind a portset to.

`-portset <text>` - Portset Binding Igroup

Specifies the portset name that you want to bind an initiator group to.

Examples

```
cluster1::> lun igroup bind -vserver vs1 -igroup ig1 -portset-name ps1
```

Binds igroup ig1 to portset ps1.

Related Links

- lun igroup unbind

lun igroup create

Create a new initiator group

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

Description

This command creates a new initiator group (igroup). Use igroups to control which hosts have access to specific LUNs. When you bind an igroup to a portset, a host in the igroup can access the LUNs only by connecting to the target ports in the portset.

When you create an igroup, you can add multiple existing initiators by specifying them in a list, separating them with commas. Later, you can add or remove initiators from the initiator group. Use the lun igroup add command to add initiators. Use the lun igroup remove command to remove an initiator. Unless the `-initiator` option is supplied, no initiators are added to a new igroup.

You can also bind a portset to an initiator when you create an initiator group. You can modify the portset binding of an initiator group by using the lun igroup bind command or the lun igroup unbind command.

The name you assign to an igroup is independent of the name of the host that is used by the host operating system, host files, or Domain Name Service (DNS). If you name an igroup aix1, for example, it is not mapped to the actual IP host name (DNS name) of the host.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the name of the new initiator group. An initiator group name is a case-sensitive name that must contain one to 96 characters. Spaces are not allowed.

> It might be useful to provide meaningful names for igroups, ones that describe the hosts that can access the LUNs mapped to them.

{ `[-protocol {mixed|fcp|iscsi}]` - Protocol

Specifies if the initiator group protocol is *`fcp`* , *`iscsi`* , or *`mixed`* .

| `[-f, -fcp <true>]` - FCP

Specifies FCP as the protocol type of the new initiator group.

| `[-i, -iscsi <true>]` - iSCSI }

Specifies iSCSI as the protocol type of the new initiator group.

`-t, -ostype <Initiator Group OS Type>` - OS Type

Specifies the operating system type for the new initiator group. The operating system type indicates the type of host operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same operating system type. The operating system types of initiators are:

- solaris

- windows
- hpux
- aix
- linux
- netware
- vmware
- openvms
- xen
- hyper_v

[-a, -portset <text>] - Portset Binding Igroup

Specifies that a portset is bound to the initiator.

{ -initiator <text>,… - Initiators

Specifies the initiators that are attached to the new initiator group. By default, no initiators are added to the new igroup.

| [-child-igroups <text>,…] - Child Igroups }

Specifies the child initiator groups that are attached to the new initiator group. The initiator group inherits the initiators of any child igroups. An initiator group can have either initiators or child initiator groups, but not both. Initiator groups can be nested to match host cluster hierarchies. If the host cluster then changes, initiators only need to be updated in the child igroups and the parent igroups inherit the changes.

[-delete-on-unmap {true|false}] - Delete on Last Unmap

Specifies if this initiator group will be deleted automatically when no longer a member of a LUN mapping relationship.

[-comment <text>] - Igroup Comment

A description for the igroup to create. If the comment string contains white space, it must be enclosed in double quotes. The limit is 254 characters.

[-replication-peer <text>] - Replication Peer

Specifies the local name of a peer Vserver to which the initiator group should be replicated. By default, the initiator group is not replicated.

[-proximal-vserver {local|peer}] - Proximal Vserver (privilege: advanced)

Specifies a Vserver to add to any initiators being added to an initiator group. *local* refers to the initiator group's Vserver. *peer* refers to the initiator group's replication peer. *proximal-vserver* can only be set for replicated initiator groups.

Examples

```
cluster1::> lun igroup create -vserver vs1 -igroup ig1 -protocol mixed
-ostype linux -initiator iqn.2001-04.com.example:abc123
```

Creates initiator group *ig1* on Vserver *vs1* with a *mixed* protocol type on a Linux operating system with the

initiator `iqn.2001-04.com.example:abc123`.

Related Links

- lun igroup add
- lun igroup remove
- lun igroup bind
- lun igroup unbind

lun igroup delete

Delete an initiator group

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

Description

This command deletes an existing initiator group. By default, you cannot delete an initiator group if LUN maps for that initiator group exist. You need to unmap all the LUNs that are associated with that initiator group before you can delete the initiator group. Use the lun unmap command to remove LUNS from an initiator group.

You can specify the `force` option to delete an initiator group and remove existing LUN maps defined for that initiator group.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the initiator group that you want to delete.

`[-f, -force <true>]` - Force

Deletes an initiator group and all associated LUN maps.

Examples

```
cluster1::> lun igroup delete -vserver vs1 -igroup ig1
```

Deletes the initiator group ig1 on Vserver vs1.

Related Links

- lun unmap

lun igroup disable-aix-support

Disables SAN AIX support on the cluster

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

Description

This command disables the SAN AIX support across the cluster (all Vservers and all AIX initiator groups). However, before you can disable SAN AIX support, you must remove all SAN AIX related objects from the cluster. You need to unmap all the LUNs that are associated with the AIX initiator groups. Then you need to delete all of the AIX initiator groups. Use the lun unmap command to remove LUNS from an initiator group. Use the `igroup delete` command to delete an initiator group.

> This command is not intended to be used in normal operation. Use only when you are downgrading to a release that does not support SAN AIX operation.

Examples

```
cluster1::> lun igroup disable-aix-support
```

Disables the SAN AIX support for cluster1.

Related Links

- lun unmap

lun igroup modify

Modify an existing initiator group

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

Description

This command modifies an attribute for an initiator group. Currently, the only settable attribute is the operating system.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the initiator group whose attribute you want to modify.

`[-t, -ostype <Initiator Group OS Type>]` - OS Type

Specifies the operating system that you want to modify. The operating system types of initiators are:

- solaris
- windows
- hpux
- aix

- linux
- netware
- vmware
- openvms
- xen
- hyper_v

[-delete-on-unmap {true|false}] - Delete on Last Unmap

Specifies if this initiator group will be deleted automatically when no longer a member of a LUN mapping relationship.

[-comment <text>] - Igroup Comment

Specifies the comment for the igroup you want to modify. If the comment string contains white space, it must be enclosed in double quotes. The limit is 254 characters. To clear the comment, specify the empty string enclosed in double quotes.

[-replication-peer <text>] - Replication Peer

Specifies the local name of a peer Vserver to which the initiator group should be replicated. To terminate replication of the initiator group, set this to "-".

Examples

```
cluster1::> lun igroup modify -vserver vs1 -igroup ig1 -ostype windows
```

Changes the operating system to *windows* for initiator group *ig1* on Vserver *vs1* .

lun igroup remove

Remove initiators from an initiator group

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

Description

This command removes an initiator from an initiator group. You can only remove an initiator if no existing LUN maps are defined for that initiator group. You must unmap the LUNs from the initiator group with the lun unmap command before you can remove initiators from the initiator group.

You can use the `force` option to remove an initiator from a mapped initiator group. Removing an initiator from a mapped initiator group makes the LUNs to which the initiator group is mapped no longer available to the initiator. This might cause a disruption in the availability of data.

Parameters

-vserver <Vserver Name> - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the initiator group from which to remove an initiator.

`[-initiator <text>,…]` - Initiators

Specifies the initiator name to remove. Use the WWPN, IQN or the alias of the initiator.

`[-f, -force <true>]` - Force

Forcibly removes an initiator from a mapped initiator group. This might cause a disruption in the availability of data.

`[-child-igroups <text>,…]` - Child Igroups

Specifies the child initiator groups to remove.

Examples

```
cluster1::> lun igroup remove -vserver vs1 -igroup ig1 -initiator
iqn.1992-08.com.mv.mvinitiator
```

Removes the initiator iqn.1992-08.com.mv.mvinitiator from Vserver vs1 and initiator group ig1.

Related Links

- lun unmap

lun igroup rename

Rename an existing initiator group

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

Description

This command renames an existing initiator group. When you rename an initiator group, this action does not affect access to the LUNs mapped to the initiator group you want to rename.

An initiator group name is a case-sensitive name and must meet the following requirements:

- Must contain one to 96 characters. Spaces are not allowed.
- Can contain the letters A through Z, a through z, numbers 0 through 9, hyphen (-), underscore (_), colon (:), and period (.).
- Must start with a letter or number.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

-igroup <text> - Igroup Name

Specifies the initiator group you want to rename.

-new-name <text> - New Igroup Name

Specifies the new name of the initiator group.

Examples

```
cluster1::> lun igroup rename -vserver vs1 -igroup ig1 -new-name ignew1
```

Renames an initiator group from ig1 to ignew1 on Vserver vs1.

lun igroup show

Display a list of initiator groups

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

Description

This command displays status information for initiator groups (igroup). By default, the command displays status for all initiator groups.

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.

[-vserver <Vserver Name>] - Vserver Name

Specifies the Vserver.

[-igroup <text>] - Igroup Name

Selects the initiator groups that match this parameter value.

[-protocol {mixed|fcp|iscsi}] - Protocol

Selects the initiator groups that match this parameter value (`FCP` , `iSCSI` , or `mixed`).

[-t, -ostype <Initiator Group OS Type>] - OS Type

Selects the initiator groups that match this parameter value. The operating system types are

- `solaris`
- `windows`
- `hpux`

- aix
- linux
- netware
- vmware
- openvms
- xen
- hyper_v

[-a, -portset <text>] - Portset Binding Igroup

Selects the initiator groups that match this parameter value.

[-initiator <text>,…] - Initiators

Selects the initiator groups that match this parameter value.

[-child-igroups <text>,…] - Child Igroups

Selects the initiator groups that match this parameter value.

[-uuid <UUID>] - Igroup UUID

Selects the initiator groups that match this parameter value.

[-delete-on-unmap {true|false}] - Delete on Last Unmap

Selects the initiator groups that match this parameter value. A value of `true` displays all the initiator groups that will be deleted automatically when they are no longer used in a LUN mapping relationship.

[-comment <text>] - Igroup Comment

Selects the igroups that match this parameter value.

[-replication-peer <text>] - Replication Peer

Selects the initiator groups that match this parameter value. The replication-peer is the local name of the peer Vserver, as reported by the vserver peer show -fields peer-vserver command. When set, the igroup and any lun maps to LUNs in SMBC relationships are replicated to the peer cluster.

[-replication-error <text>] - Replication Error

Selects the initiator groups that match this parameter value. This field provides information about asynchronous errors encountered while replicating this igroup. Igroups within a peering relationship are replicated in the same stream, so the error reported here might be related to this igroup or a prior replicated igroup that is now blocking the replication of this igroup. Both the error information and the igroup encountering the error are reported. If the error is configuration related, it can be corrected on the referenced igroup. The replication is retried using exponential backoff up to a maximum of one retry every 5 minutes. Every operation on the same stream triggers an immediate retry and restarts the exponential backoff starting with a 1 second delay. If the error is system related, the retries should correct the error when the system enters a healthy state.

[-replication-error-igroup <text>] - Replication Error Igroup

Selects the initiator groups that match this parameter value. The name of the igroup encountering the error reported by the 'replication-error' field.

`[-replicated-igroup-unreplicated-luns {true|false}]` - Replicated Igroup Contains Unreplicated LUNs

Selects the initiator groups that match this parameter value. This field indicates that the initiator group is replicated but contains one or more unreplicated LUNs.

Examples

```
cluster1::> igroup show -instance
          Vserver Name: vs0
           Igroup Name: ig1
              Protocol: mixed
               OS Type: linux
Portset Binding Igroup: -
           Igroup UUID: 358338ba-cfd6-11df-a9ab-123478563412
            Initiators: iqn.1992-08.com.mv:abc (not logged in)
Vserver Name: vs0
           Igroup Name: ig2
              Protocol: mixed
               OS Type: linux
Portset Binding Igroup: -
           Igroup UUID: 3fb136c7-cfd6-11df-a9ab-123478563412
            Initiators: -
Vserver Name: vs1
           Igroup Name: ig1
              Protocol: mixed
               OS Type: windows
Portset Binding Igroup: p1
           Igroup UUID: 03accf6b-d08c-11df-a9ab-123478563412
            Initiators: -
3 entries were displayed.
```

The example above displays information about all initiator groups.

Related Links

- vserver peer show

lun igroup unbind

Unbind an existing initiator group from a portset

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

Description

This command unbinds an initiator group from a portset. When you unbind an initiator group from a portset, all of the initiators in the initiator group have access to all target LUNs on all network interfaces.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the initiator group that you want to unbind from the portset.

Examples

```
cluster1::> lun igroup unbind -vserver vs1 -igroup ig1
```

Unbinds the initiator group ig1 from the portset on Vserver vs1.

lun igroup initiator add-proximal-vserver

Add a Vserver to the initiator's proximal list

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

Description

This command is used to configure which site in a SnapMirror Active Synchronous (SM-AS) relationship is proximal with the provided FCP or iSCSI initiator.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Specifies the name of the Vserver the initiator will access.

`-initiator <text>` - Initiator (privilege: advanced)

Specifies the FCP WWPN or iSCSI IQN of the initiator.

`-proximal-vservers <text>,…` - Proximal Vserver Name (privilege: advanced)

Specifies the name of the Vserver to configure as proximal. The specified Vserver must either be the same as the `-vserver` parameter or a Vserver in a peer relationship with *`snapmirror`* as an allowed application.

`[-replace <true>]` - Replace Existing (privilege: advanced)

If true, replace all existing `-proximal-vservers` for the initiator with the ones specified. By default, the specified `-proximal-vservers` are added to any existing values.

Examples

The following example configures *`iqn.2001-04.com.example:uefi-0232314`* in Vserver *`local-vserver`* as proximal to the *`peer-vserver`* Vserver:

```
cluster1::*> lun igroup initiator add-proximal-vserver -vserver local-
vserver -initiator iqn.2001-04.com.example:uefi-0232314 -proximal-vservers
peer-vserver
```

lun igroup initiator modify

Modify an initiator

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

Description

This command modifies an attribute for an initiator.

Parameters

`-vserver <Vserver Name>` - Vserver

Specifies the vserver of the initiator you want to modify.

`-initiator <text>` - Initiator

Specifies the initiator whose attribute you want to modify.

`[-comment <text>]` - Comment

The initiator's new comment. The comment has a maximum length of 255 characters. To clear the current comment, provide an empty string.

Examples

```
cluster1::> lun igroup initiator modify -vserver vs1 -initiator iqn.1991-
05.com.example:name -comment "My application host"
```

Changes the comment to "My application host" for initiator *`iqn.1991-05.com.example:name`* in Vserver *`vs1`* .

```
cluster1::> lun igroup initiator modify -vserver vs1 -initiator iqn.1991-
05.com.example:name -comment ""
```

Clears the comment for initiator *`iqn.1991-05.com.example:name`* in Vserver *`vs1`* .

lun igroup initiator remove-proximal-vserver

Remove a Vserver from the initiator's proximal list

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

Description

This command is used to remove the proximal Vserver configuration from an FCP or iSCSI initiator.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Specifies the name of the Vserver the initiator will access.

`-initiator <text>` - Initiator (privilege: advanced)

Specifies the FCP WWPN or iSCSI IQN of the initiator.

{ `-proximal-vservers <text>,…` - Proximal Vserver Name (privilege: advanced)

Specifies the name of the Vserver to remove as proximal. The specified Vserver must either be the same as the `-vserver` parameter or a Vserver in a peer relationship with *`snapmirror`* as an allowed application.

| `-all <true>` - Remove All (privilege: advanced) }

If true, remove all existing proximal Vservers from the specified initiator.

Examples

The following example configures *`iqn.2001-04.com.example:uefi-0232314`* in Vserver *`local-vserver`* as not proximal to the *`peer-vserver`* Vserver:

```
cluster1::*> lun igroup initiator remove-proximal-vserver -vserver local-
vserver -initiator iqn.2001-04.com.example:uefi-0232314 -proximal-vservers
peer-vserver
```

lun igroup initiator show

Display initiators

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

Description

This command displays information for initiators that are members of initiator groups (igroups).

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.

[-vserver <Vserver Name>] - Vserver

Selects the initiators of igroups whose Vserver matches the specified value.

[-initiator <text>] - Initiator

Selects initiators whose name matches the specified value.

[-comment <text>] - Comment

Selects initiators whose comment matches the specified value.

[-igroups <text>,…] - Igroups

Selects initiators belonging to the specified igroups.

[-proximal-vservers <text>,…] - Proximal Vservers (privilege: advanced)

Selects initiators configured as proximal to the specified Vservers.

Examples

```
cluster1::> lun igroup initiator show -instance
  Vserver: vs0
Initiator: init1
  Comment: Initiator Comment
  Igroups: ig1,ig2
Vserver: vs0
Initiator: init2
  Comment: Initiator Comment2
  Igroups: ig3
2 entries were displayed.
```

This example displays detailed information for all initiators in igroups.

lun import commands

lun import create

Create an import relationship

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

Description

This command creates an import relationship between a specified LUN and a specified foreign disk so you can import the foreign disk data into a LUN.

The foreign disk must be marked as foreign using storage disk set-foreign-lun command before you can begin the import progress.

The LUN must be of the same size as the foreign disk.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN where you import data to from the foreign disk data.

`-foreign-disk <text>` - Foreign Disk Serial Number (privilege: advanced)

Specifies the serial number of the Foreign Disk.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN where you want to import the data of the foreign disk to. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

Examples

```
cluster1::> lun import create -vserver vs1 -path /vol/dvol1/lun1 -foreign
-disk 6000B5D0006A0000006A020E00040000
```

Creates an import relationship between lun1 at the path /vol/dvol1/lun1 and foreign disk with serial number 6000B5D0006A0000006A020E00040000.

Related Links

- storage disk set-foreign-lun

lun import delete

Deletes the import relationship of the specified LUN or the specified foreign disk

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

Description

This command deletes the import relationship of a specified LUN or a specified foreign disk.

You cannot use this command if an import is in-progress between the foreign disk and the LUN unless you use the force option. The import has to either successfully completed or be stopped before deleting the import relationship.

You can use the lun import stop command to stop the data import, and then you delete the import relationship.

Parameters

{ `-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN that you want to delete the import relationship.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN where you want to delete the import relationship. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

| -foreign-disk <text> - Foreign Disk Serial Number (privilege: advanced) }

Specifies the serial number of the foreign disk.

[-force {true|false}] - Force Delete (privilege: advanced)

When set to true, stops the in progress data import.

Examples

```
cluster1::> lun import delete -vserver vs1 -path /vol/vol2/lun2
```

Deletes the import relationship of lun2 at the path /vol/vol2/lun2.

```
cluster1::> lun import delete -vserver vs0 -foreign-disk
6000B5D0006A0000006A020E00040000
```

Deletes the import relationship of the foreign disk with serial number 6000B5D0006A0000006A020E00040000.

Related Links

- lun import stop

lun import pause

Pause the import for the specified LUN

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

Description

This command pauses the data import to a specified LUN.

This command does not reset all import checkpoints. To resume a paused import, use the lun import resume command to retart from the last checkpoint taken before you paused the data import.

If you want to resume the data import from the beginning, use the lun import stop command. Then use the lun import start command.

Parameters

-vserver <Vserver Name> - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN you want to pause the data import to.

-path <path> - LUN Path (privilege: advanced)

Specifies the path of the LUN you want to pause the data import to. Examples of correct LUN paths are */vol/vol1/lun1* and */vol/vol1/qtree1/lun1* .

Examples

```
cluster1::> lun import pause -vserver vs1 -path /vol/vol2/lun2
```

Pauses the data import for lun2 at the path /vol/vol2/lun2

lun import prepare-to-downgrade

Prepares LUN import to be downgraded

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

Description

This command prepares the cluster for a downgrade to a version of ONTAP earlier than 8.3.1 by disabling the online LUN import feature. Before using this command verify that all LUNs in an import relationships are offline by running lun show .

Examples

```
cluster1::> lun import prepare-to-downgrade
```

The above example will disable the online LUN import feature if all LUNs in import relationships are offline

Related Links

- lun show

lun import resume

Resume the import for the specified LUN

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

Description

Resumes the data import to a specified LUN.

The import starts from the last checkpoint taken before you paused the data import.

If you want to resume the data import from the beginning, use the lun import stop command. Then use the lun import start command.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN you want to resume the data import to.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN that you want to resume the data import to. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

Examples

```
cluster1::> lun import resume -vserver vs1 -path /vol/vol2/lun2
```

Resumes the data import to lun2 at the path /vol/vol2/lun2

lun import show

Display a list of import relationships

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

Description

This command displays information about the import relationships.

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.

`[-vserver <Vserver Name>]` - Vserver Name (privilege: advanced)

Displays import relationships for a specified Vserver.

`[-foreign-disk <text>]` - Foreign Disk Serial Number (privilege: advanced)

Enables you to see the import relationship for a particular foreign disk with the specified serial number.

`[-path <path>]` - LUN Path (privilege: advanced)

Enables you to see the import relationship for a particular LUN path. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

`[-import-home-node {<nodename>|local}]` - Import Home Node (privilege: advanced)

Enables you to see the node that initially started the data import and where the I/O for the foreign disk is directed. If failover occurs, any in-progress data import restarts on the partner node.

`[-import-current-node {<nodename>|local}]` - Import Current Node (privilege: advanced)

Displays the node that is currently importing data and where the I/O for the foreign disk is directed. During giveback and if the import home node is different from the current home node , import restarts on the initial node (import-home-node).

[-operation-in-progress {import|verify}] - Operation in Progress (privilege: advanced)

Enables you to see the imports in progress or import verification in progress.

[-admin-state {stopped|started|paused}] - Admin State (privilege: advanced)

Enables you to see the import relationships for a specified administrative state. For example, you can list all the imports that have started in a cluster.

[-operational-state {in_progress|failed|completed|stopped|paused}] - Operational State (privilege: advanced)

Enables you to see the import relationships for a specified operational state. For example, you can list all the imports that have completed in a cluster.

[-percent-complete <integer>] - Percent Complete (privilege: advanced)

Enables you to see the percentage of completion for both import and verification. If you want to see all the complete imports and verifications, you would specify 100.

[-imported-blocks <integer>] - Blocks Imported (privilege: advanced)

Enables you to see the number of blocks that have been imported to the LUN.

[-compared-blocks <integer>] - Blocks Compared (privilege: advanced)

Enables you to see the number of LUN and foreign disk blocks that have been compared.

[-total-blocks <integer>] - Total Blocks (privilege: advanced)

Enables you to see the total number of blocks that must be imported to complete the data import to a LUN or the number of blocks that must be compared to complete the import verification.

[-estimated-remaining-duration {<seconds>|[<d> days] <hh>:<mm>[:<ss>]}] - Estimated Remaining Duration (privilege: advanced)

If this parameter is specified, the command displays import or verify operations that match the specified time.

[-failure-reason <text>] - Failure Reason (privilege: advanced)

Selects LUN import operations that match this parameter value.

[-max-throughput-limit {<integer>[KB|MB|GB|TB|PB]}] - Maximum Throughput Limit (per sec) (privilege: advanced)

Selects the LUN import operations that match this parameter value. This value is the throughput limit at which an import or verify will be throttled. By default, there is no throttling.

[-current-throughput {<integer>[KB|MB|GB|TB|PB]}] - Current Throughput (per sec) (privilege: advanced)

Selects the LUN import operations that match this parameter value. This value is the current throughput for an in-progress import or verify operation.

[-qos-policy-group <text>] - QoS Policy Group (privilege: advanced)

Selects the LUN import operations that match this parameter value. This value is the QoS policy group associated with an import relationship.

Examples

```
cluster1::> lun import show
vserver           foreign-disk                      path
operation-in-progress   admin-state  operational-state  percent-complete
 ---------------------------------------------------------------------------
---------------------------------------------------------------------------
-
vs1               6000B5D0006A0000006A020E00040000  /vol/dvol1/lun1
import                  stopped      stopped            0
vs1               60060480343631336433336538366537  /vol/vol1/lun1
import                  started      failed            11
vs2               6000B5D0006A0000006A020E00040001  /vol/dvol1/lun2
verify                  started      in_progress        5
```

Display information about all import relationships in the cluster

```
cluster1::> lun import show -instance
                      Vserver Name: vs1
                          LUN Path: /vol/dvol1/lun1
        Foreign Disk Serial Number: 6000B5D0006A0000006A020E00040000
                  Import Home Node: cluster1-01
               Current Import Node: cluster1-01
             Operation in Progress: import
                       Admin State: started
                Operational State: in-progress
                  Percent Complete: 0%
                   Blocks Imported: 0
                   Blocks Compared: 0
            Total Blocks to Import: 10000000
      Estimated Remaining Duration: 00:01:23
                     Failure Reason: -
Maximum Throughput Limit (per sec): -
      Current Throughput (per sec): -
                  QoS Policy Group: -
Vserver Name: vs2
                          LUN Path: /vol/dvol1/lun2
        Foreign Disk Serial Number: 6000B5D0006A0000006A020E00040001
                  Import Home Node: cluster1-01
               Current Import Node: cluster1-01
             Operation in Progress: verify
                       Admin State: started
                Operational State: in-progress
                  Percent Complete: 5%
                   Blocks Imported: 10000000
```

```
                    Blocks Compared: 500000
            Total Blocks to Import: 10000000
      Estimated Remaining Duration: 00:00:59
                     Failure Reason: -
Maximum Throughput Limit (per sec): 2MB
      Current Throughput (per sec): 1.29MB
                  QoS Policy Group: fli_pg_cf2b638b-606b-11e4-ae4c-
000c290d40ff
Vserver Name: vs1
        Foreign Disk Serial Number: 600604803436313364333365383665 37
                          LUN Path: /vol/vol1/lun1
                  Import Home Node: cluster1-01
               Current Import Node: cluster1-01
             Operation in Progress: import
                       Admin State: started
                 Operational State: failed
                  Percent Complete: 11
                   Blocks Imported: 932352
                   Blocks Compared: -
                      Total Blocks: 8388608
      Estimated Remaining Duration: -
                     Failure Reason: Source read error - reservation
conflict.
Maximum Throughput Limit (per sec): 12MB
      Current Throughput (per sec): -
                  QoS Policy Group: fli_pg_f6632344-60e7-11e4-9bad-
000c290d40ff
```

Display detailed information about all import relationships in the cluster.

```
cluster1::> lun import show -vserver vs1
vserver   path                  foreign-disk                      admin-
state  operational-state  percent-complete
 ---------------------------------------------------------------------------
----------------------------------------
vs1       /vol/dvol1/lun1       vgv3040f46a:vgbr300s70:9.126L1    stop
-                 0%
```

Display information about the LUNs in an import relationships in a specific vserver.

```
cluster1::> lun import show -admin-state start
vserver   path                  foreign-disk                      admin-
state  operational-state  percent-complete
 ---------------------------------------------------------------------------
----------------------------------------
vs2       /vol/dvol1/lun2       vgv3040f46a:vgbr300s70:9.126L2    start
in-progress        5%
```

Display active LUN import sessions in a cluster.

lun import start

Start the import for the specified LUN

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

Description

This command initiates the data import to a specified LUN.

You must use the lun import create command to create an import relationship between a LUN and a foreign disk before you can initiate the data import.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN you want to import data to.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN that you want to import data to. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

Examples

```
cluster1::> lun import start -vserver vs1 -path /vol/vol2/lun2
```

Starts the data import to lun2 at the path /vol/vol2/lun2.

lun import stop

Stop the import for the specified LUN

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

Description

This command stops the data import into a specified LUN.

After you stop the data import and if you start the import again using lun import start command, then the import restarts from the beginning.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN you want to stop importing data to.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN that you want to stop the data import to.

Examples

```
cluster1::> lun import stop -vserver vs1 -path /vol/vol2/lun2
```

Stops data import to lun2 at the path /vol/vol2/lun2

Related Links

- lun import start

lun import throttle

Modify the max throughput limit for the specified import relationship

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

Description

This command throttles the speed of the import for a given LUN by specifying a maximum throughput limit on the import.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN to which data from the foreign disk is imported.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN to which data from the foreign disk is imported. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

`-max-throughput-limit {<integer>[KB|MB|GB|TB|PB]}` - Maximum Throughput Limit (per sec) (privilege: advanced)

Specifies the maximum amount of throughput to be allocated for processing import requests on the bound LUN. At the time of creation, the default is zero. A value of zero implies that import traffic is processed by the system at best effort rate along with on-going user I/O. A non-zero value indicates that import will be throttled at a rate which is at most the maximum throughput limit set.

Examples

```
cluster1::*> lun import throttle -vserver vs1 -path /vol/vol1/lun1 -max
-throughput-limit 3M
```

The above example limits the import speed for the bound LUN with path /vol/vol1/lun1 to a maximum of 3MB/s.

lun import verify start

Start the verification of the foreign disk and LUN data

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

Description

This command compares the LUN and the foreign disk block by block. You are not required to run this command; it is optional. Before you can do this verification process, the operation state must be stopped or completed. Use the lun import show command to determine the operation state.

If a block mismatch occurs, the verification process stops.

Verification must be done offline. Ensure the foreign disk and LUN cannot be accessed by a host. To prevent access of the LUN, the LUN should be taken offline administratively using the lun offline command.

Note: The specified LUN must be in an import relationship with a foreign disk before you can verify the data import.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN you want to compare block by block with the foreign disk.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN that you want to compare the foreign disk to. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

Examples

```
cluster1::> lun import verify start -vserver vs1 -path /vol/vol2/lun2
```

Starts the import verification on lun2 at path /vol/vol2/lun2.

lun import verify stop

Stop the verify for the specified LUN

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

Description

This command stops the block by block verification of the foreign disk and LUN data.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver that contains the LUN you want to stop block by block comparison with the foreign disk.

`-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN that you want to stop the block by block comparison. Examples of correct LUN paths are *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`* .

Examples

```
cluster1::> lun import verify stop -vserver vs1 -path /vol/vol2/lun2
```

Stops the import verify on lun at path /vol/vol2/lun2.

lun mapping commands

lun mapping add-reporting-nodes

Add Reporting Nodes

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

Description

This command is used before or after a data mobility event that modifies the owning node of the LUN to add the new optimized nodes to the specified LUN mapping's reporting nodes.

For more information on managing reporting nodes in response to data mobility events, see the ONTAP SAN Administration Guide.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the name of the Vserver containing the LUN.

{ `-path <path>` - LUN Path

Specifies the path of the LUN. Examples of correct LUN paths are: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name

Specifies the volume that contains the LUN.

[-qtree <qtree name>] - Qtree Name

Specifies the qtree that contains the LUN.

-lun <text> - LUN Name }

Specifies the LUN name.

-g, -igroup <text> - Igroup Name

Specifies the igroup the LUN is mapped to.

{ -local-nodes <true> - Add Nodes for Current LUN Location

Add the current LUN owner node and HA partner to the LUN mapping's reporting nodes.

This option should be used after a LUN mobility event to restore optimized access to the LUN.

| -destination-aggregate <aggregate name> - Add Nodes for Aggregate

Add the specified aggregate's owner node and HA partner to the LUN mapping's reporting nodes.

This option may be used prior to a LUN mobility event that changes the LUN's containing aggregate.

| -destination-volume <volume name> - Add Nodes for Volume }

Add the specified volume's owner node and HA partner to the LUN mapping's reporting nodes.

This option may be used prior to a LUN mobility event that changes the LUN's containing volume.

| -all <true> - Add All Nodes (privilege: advanced) }

Set the LUN mapping to report on all nodes in preparation for a revert to a previous version of ONTAP.

Examples

```
cluster1::> lun mapping add-reporting-nodes -vserver vs1 -path
/vol/vol1/lun1 -igroup ig1
```

Add the current owner node and HA partner for the LUN mapping of */vol/vol1/lun1* to igroup *ig1* .

```
cluster1::> lun mapping add-reporting-nodes -vserver vs1 -lun * -igroup
ig1 -destination-aggregate aggr2
```

Add the aggregate owner node and HA partner for aggregate *aggr2* to all LUN mappings in volume *vol1* to igroup *ig1* prior to starting a volume move operation.

lun mapping create

Map a LUN to an initiator group

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

Description

This command maps a LUN to all of the initiators in an initiator group (igroup). After you map the LUN, the LUN is visible to all initiators in the igroup.

ONTAP ensures that there are no LUN map conflicts whether the LUN is offline or online. A LUN map conflict is a mapping that would violate either of the following rules:

- Each LUN can be mapped to an initiator only once. A LUN can be mapped to multiple igroups as long as each igroup has a distinct set of initiators.
- LUN IDs must be unique such that every initiator has a unique ID for each LUN to which it is mapped. If you map a LUN to an igroup, the LUN ID for that mapping cannot be reused by any of the initiators in that igroup.

In order to determine if a LUN ID is valid for a mapping, ONTAP checks each initiator in the igroup to make sure that the LUN ID is not used for another mapping that includes that initiator.

> Prior to mapping a LUN, you must have at least one iSCSI or FCP LIF provisioned on the LUN's owner node and high-availability partner node.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver that contains the LUN you want to map.

{ `-path <path>` - LUN Path

Specifies the path of the LUN that you want to map. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name

Specifies the volume that contains the LUN you want to map.

`[-qtree <qtree name>]` - Qtree Name

Specifies the qtree that contains the LUN you want to map.

`-lun <text>` - LUN Name }

Specifies the LUN name that you want to map.

`-g, -igroup <text>` - Igroup Name

Specifies the igroup that you want to map.

`[-lun-id <integer>]` - LUN ID

Specifies the LUN ID for the mapping. The LUN ID is specific to the mapping, not to the LUN itself. This is used by the initiators in the igroup as the Logical Unit Number for the initiator when accessing the storage.

`[-additional-reporting-node <nodename>]` - Additional Reporting Node (privilege: advanced)

Specifies an additional node to populate the `-reporting-nodes` list when creating the LUN mapping. The specified node's high availability partner will be automatically populated as well. Use this parameter when preferred data mobility destinations are known ahead of time and the appropriate paths can be pre-configured.

Examples

```
cluster1::> lun mapping create -vserver vs1 -path /vol/vol1/lun1 -igroup
ig1 -lun-id 8
```

Maps a LUN at /vol/vol1/lun1 on Vserver vs1 to the igroup ig1 with LUN ID 8.

lun mapping delete

Unmap a LUN from an initiator group

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

Description

This command unmaps a LUN from an initiator group. After you use this command, the LUN is not visible to any of the initiators in the initiator group.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Selects the LUN maps for the Vserver that matches the parameter value.

{ `-path <path>` - LUN Path

Specifies the path of the LUN you want to unmap. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name

Specifies the volume of the LUN you want to unmap.

`-qtree <qtree name>` - Qtree Name

Specifies the qtree of the LUN you want to unmap.

`-lun <text>` - LUN Name }

Specifies the name of the LUN you want to unmap.

`-g, -igroup <text>` - Igroup Name

Specifies the initiator group that you want to unmap the LUN from.

Examples

```
cluster1::> lun mapping delete -vserver vs1 -path /vol/vol1/lun1 -igroup
ig1
```

Unmaps LUN at path /vol/vol1/lun1 from the initiator group ig1 on Vserver vs1.

lun mapping remove-reporting-nodes

Remove Reporting Nodes

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

Description

This command is used after a data mobility event to remove reporting nodes that are no longer required for optimized access from the specified LUN mapping.

For more information on managing reporting nodes in response to data mobility events, see the ONTAP SAN Administration Guide.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the name of the Vserver containing the LUN.

{ `-path <path>` - LUN Path

Specifies the path of the LUN. Examples of correct LUN paths are: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name

Specifies the volume that contains the LUN.

`[-qtree <qtree name>]` - Qtree Name

Specifies the qtree that contains the LUN.

`-lun <text>` - LUN Name }

Specifies the LUN name.

`-g, -igroup <text>` - Igroup Name

Specifies the igroup the LUN is mapped to.

`-remote-nodes <true>` - Remove Remote Nodes for LUN Location

If specified, remove all nodes other than the LUN's owner and HA partner from the LUN mapping's reporting nodes.

Examples

```
cluster1::> lun mapping remove-reporting-nodes -vserver vs1 -path
/vol/vol1/lun1 -igroup ig1
```

Remove excess remote nodes from the LUN mapping of *`/vol/vol1/lun1`* to igroup *`ig1`*

```
cluster1::> lun mapping remove-reporting-nodes -vserver vs1 -path
/vol/vol1/lun1 -igroup ig1
```

Remove excess remote nodes from the LUN mapping of `/vol/vol1/lun1` to igroup `ig1` .

lun mapping show-initiator

Show the LUN mappings to a specific initiator

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

Description

The `lun mapping show-initiator` command lists the LUNs which are mapped to an initiator group which contains a specific initiator.

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.

`[-vserver <Vserver Name>]` - Vserver Name (privilege: advanced)

Selects the LUN mappings for the vserver that you specify.

`-initiator <text>` - Initiator Name (privilege: advanced)

Selects the LUN mappings for the initiator that you specify.

`[-lun-id <integer>]` - Logical Unit Number (privilege: advanced)

Selects the LUN mappings with a LUN ID that you specify.

`[-igroup <text>]` - Igroup Name (privilege: advanced)

Selects the LUN mappings for the initiator group that you specify.

`[-path <path>]` - LUN Path (privilege: advanced)

Selects the LUN mappings for the LUN path that you specify.

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

Selects the LUN mappings for the LUNs which are being hosted on the node that you specify.

`[-reporting-nodes <nodename>,…]` - Reporting Nodes (privilege: advanced)

Selects the LUN mappings for the LUNs which have reporting nodes that you specify.

`[-vserver-uuid <UUID>]` - Vserver UUID (privilege: advanced)

Selects the LUN mappings for the Vserver UUID that you specify.

`[-igroup-uuid <UUID>]` - Igroup UUID (privilege: advanced)

Selects the LUN mappings for the initiator group UUID that you specify.

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

Selects the LUN mappings for the LUN UUID that you specify.

Examples

The following example displays the LUN mappings for initiator *20:10:0a:50:00:01:01:01* in Vserver *vs1* .

```
cluster1::> lun mapping show-initiator -vserver vs1 -initiator
20:10:0a:50:00:01:01:01
Vserver Initiator LUN ID Path                                IGroup
------- --------- ------ ------------------------------------
----------------
vs1     20:10:0a:50:00:01:01:01
                       0 /vol/igroup_1_1_vol/lun1            igroup_1
                       2 /vol/igroup_1_1_vol/lun3            igroup_1
                       3 /vol/igroup_1_2_vol/lun1            igroup_1
                       5 /vol/igroup_1_2_vol/lun3            igroup_1
                       6 /vol/igroup_1_3_vol/lun1            igroup_1
                       8 /vol/igroup_1_3_vol/lun3            igroup_1
                       9 /vol/igroup_1_4_vol/lun1            igroup_1
                      11 /vol/igroup_1_4_vol/lun3            igroup_1
                      12 /vol/igroup_2_1_vol/lun1            igroup_2
                      14 /vol/igroup_2_1_vol/lun3            igroup_2
                      15 /vol/igroup_2_2_vol/lun1            igroup_2
                      17 /vol/igroup_2_2_vol/lun3            igroup_2
                      18 /vol/igroup_2_3_vol/lun1            igroup_2
                      20 /vol/igroup_2_3_vol/lun3            igroup_2
                      21 /vol/igroup_2_4_vol/lun1            igroup_2
                      23 /vol/igroup_2_4_vol/lun3            igroup_2
16 entries were displayed.
```

lun mapping show

Lists the mappings between LUNs and initiator groups.

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

Description

This command lists the mappings between LUNs and initiator groups.

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.

[-vserver <Vserver Name>] - Vserver Name

Selects the LUN maps for the Vserver that matches the parameter value.

{ [-path <path>] - LUN Path

Selects the LUN maps for the LUN with the path that matches the parameter value. Examples of correct LUN paths: */vol/vol1/lun1* and */vol/vol1/qtree1/lun1*

| [-volume <volume name>] - Volume Name

Selects the LUN maps for the volumes that match the parameter value.

[-qtree <qtree name>] - Qtree Name

Selects the LUN maps for the queue trees that match the parameter value.

[-lun <text>] - LUN Name }

Selects the LUN maps for the LUNs with a name that matches the parameter value.

[-g, -igroup <text>] - Igroup Name

Selects the LUN maps for the igroup that matches the parameter value.

[-ostype <Initiator Group OS Type>] - Igroup OS Type

Selects the LUN maps for the initiator groups with the OS type that matches the parameter value.

[-protocol <protocol_enum>] - Igroup Protocol Type

Selects the LUN maps for initiator groups with a protocol that matches the parameter value. Possible values include FCP, iSCSI, or mixed.

[-lun-id <integer>] - LUN ID

Selects the LUN maps with a LUN ID that matches the parameter value.

[-portset <text>] - Portset Binding Igroup

Selects the LUN maps for initiator groups bound to the portset that matches the parameter value.

[-alua {true|false}] - ALUA

Selects the LUN maps with ALUA settings that match the parameter value.

[-n, -initiators <text>,…] - Initiators

Selects the LUN maps for initiator groups containing the initiators that match the parameter value.

[-node <nodename>] - LUN Node

Selects the LUN maps for nodes that match the parameter value.

[-reporting-nodes <nodename>,…] - Reporting Nodes

Selects the LUN maps that match the parameter value.

Examples

Displays all of the mappings between LUNs and initiator groups and the LUN ID for each mapping.

```
cluster1::> lun mapping show
Vserver    Path                                      Igroup   LUN ID
Protocol
---------- ----------------------------------------  -------  ------
--------
vs1        /vol/vol1/lun1                            igroup1      10
mixed
vs1        /vol/vol1/lun1                            igroup2       4
mixed
vs1        /vol/vol5/lun1                            igroup3       6
mixed
vs1        /vol/vol5/lun2                            igroup3       1
mixed
4 entries were displayed.
```

lun move commands

lun move cancel

Cancel a LUN move operation before the new LUN has been created

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

Description

The `lun move cancel` command cancels an ongoing LUN move operation prior to creation of the new LUN. The command fails if the LUN already exists at the destination path; in that case, allow the current move operation to complete and then move it back using the lun move start command.

All data transfers will be halted. If the source LUN was quiesced, it will be restored to normal operation.

> This is an advanced command because the preferred way to cancel a LUN move operation is to wait until the new LUN becomes visible, and then move it back.

Parameters

{ -vserver <Vserver Name> - Vserver Name (privilege: advanced)

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path (privilege: advanced)

Specifies the full path to the new LUN, in the format /vol/<volume>[/<qtree>]/<lun>.

Examples

```
cluster1::*> lun move cancel -vserver vs1 -destination-path /vol/vol2/lun2
```

Cancels the move operation identified by Vserver *vs1* and destination path */vol/vol2/lun2* .

Related Links

- lun move start

lun move modify

Modify an ongoing LUN move operation

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

Description

The `lun move modify` command modifies the maximum throughput of an ongoing move operation.

Parameters

{ -vserver <Vserver Name> - Vserver Name

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path

Specifies the full path to the new LUN, in the format /vol/<volume>[/<qtree>]/<lun>.

-max-throughput {<integer>[KB|MB|GB|TB|PB]} - Maximum Transfer Rate (per sec)

Specifies the maximum amount of data, in bytes, that can be transferred per second in support of this operation. This mechanism can be used to throttle a transfer, to reduce its impact on the performance of the source and destination nodes.

> The specified value will be rounded up to the nearest megabyte.

Examples

```
cluster1::> lun move modify -vserver vs1 -destination-path /vol/vol2/lun2
-max-throughput 25MB
```

Modifies the maximum throughput for the ongoing move job identified by Vserver *vs1* and destination path */vol/vol2/lun2* to 25 MB/sec.

lun move pause

Pause an ongoing LUN move operation

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

Description

The `lun move pause` command pauses an ongoing move operation. Use the lun move resume command to resume the move operation.

Parameters

{ -vserver <Vserver Name> - Vserver Name

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path

Specifies the full path to the new LUN, in the format /vol/<volume>[/<qtree>]/<lun>.

Examples

```
cluster1::> lun move pause -vserver vs1 -destination-path /vol/vol2/lun2
```

Pauses the ongoing move operation identified by Vserver *vs1* and destination path */vol/vol2/lun2* .

Related Links

- lun move resume

lun move resume

Resume a paused LUN move operation

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

Description

The `lun move resume` command resumes a paused move operation.

Parameters

{ -vserver <Vserver Name> - Vserver Name

Specifies the name of the Vserver that will host the destination LUN.

-destination-path <path> - Destination Path

Specifies the full path to the new LUN, in the format /vol/<volume>[/<qtree>]/<lun>.

Examples

```
cluster1::> lun move resume -vserver vs1 -destination-path /vol/vol2/lun2
```

Resumes the paused move operation identified by Vserver *vs1* and destination path */vol/vol2/lun2* .

lun move show

Display a list LUNs currently being moved

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

Description

The `lun move show` command shows information about LUNs currently being moved in the cluster.

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.

[-vserver <Vserver Name>] - Vserver Name

Selects LUN move operations that match this parameter value.

[-destination-path <text>] - Destination Path

Selects LUN move operations that match this parameter value.

[-source-path <path>] - Source Path

Selects LUN move operations that match this parameter value.

[-is-promoted-late {true|false}] - Is Destination Promoted Late

Selects LUN move operations that match this parameter value.

[-max-throughput {<integer>[KB|MB|GB|TB|PB]}] - Maximum Transfer Rate (per sec)

Selects LUN move operations that match this parameter value.

[-job-status {Preparing|Allocation-Map|Data|Destroying|Paused-Manual|Paused-Error|Complete|Destroyed|Terminated-Manual}] - LUN Move Status

Selects LUN move operations that match this parameter value. The possible values are:

- `Preparing` - the LUN move job is in Preparing status.
- `Allocation-Map` - the LUN move job is in Allocating status.
- `Data` - the LUN move job is in Moving Data status.
- `Destroying` - the LUN move job is in Destroying status.
- `Paused-Manual` - the LUN move job is in Manually Paused status.
- `Paused-Error` - the LUN move job is in Paused By Error status.
- `Complete` - the LUN move job is in Complete status.
- `Destroyed` - the LUN move job is in Destroyed status.

- `Terminated-Manual` - the LUN move job is in Manually Terminated status.

[-progress-percent <percent>] - LUN Move Progress (%)

Selects LUN move operations that match this parameter value.

[-elapsed-time <time_interval>] - Elapsed Time

Selects LUN move operations that match this parameter value.

[-cutover-time <time_interval>] - Cutover Time

Selects LUN move operations that match this parameter value.

[-is-snapshot-fenced {true|false}] - Is Snapshot Fenced

Selects LUN move operations that match this parameter value.

[-is-destination-ready {true|false}] - Is Destination Ready

Selects LUN move operations that match this parameter value.

[-last-failure-reason <text>] - Last Failure Reason

Selects LUN move operations that match this parameter value.

Examples

```
cluster1::> lun move show
Vserver   Destination Path                Status          Progress
--------- ------------------------------- --------------- --------
vs1       /vol/vol2/lun1                  Data            35%
vs1       /vol/vol2/lun2                  Complete        100%
2 entries were displayed.
```

The example above displays information about all the LUN move operations in the cluster.

```
cluster1::> lun move show -vserver vs1 -destination-path /vol/vol2/lun1
-instance
Destination Vserver Name: vs1
               Destination Path: /vol/vol2/lun1
                    Source Path: /vol/vol1/lun1
  Is Destination Promoted Early: false
Maximum Transfer Rate (per sec): 0B
                LUN Move Status: Data
          LUN Move Progress (%): 35%
                   Elapsed Time: 145s
            Cutover Time (secs): 0s
             Is Snapshot Fenced: true
           Is Destination Ready: true
            Last Failure Reason: -
```

The example above displays all information about the LUN being moved to `/vol/vol2/lun1` in Vserver `vs1` .

lun move start

Start moving a LUN from one volume to another within a Vserver

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

Description

The `lun move start` command initiates moving of a LUN from one volume to another. The destination volume can be located on the same node as the original volume or on a different node.

> Use lun move-in-volume command if you want to rename the LUN or move it within the same volume.

> This command does not support movement of LUNs that are created from files.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the name of the Vserver that will host the new LUN.

| `-destination-path <path>` - Destination Path

Specifies the full path to the new LUN, in the format /vol/<volume>[/<qtree>]/<lun>.

`-source-path <path>` - Source Path }

Specifies the full path to the source LUN, in the format /vol/<volume>[/<qtree>]/<lun>.

`[-promote-late <true>]` - Promote Late

Optionally specifies that the destination LUN needs to be promoted late.

If the destination is promoted early, the new LUN will be visible immediately. However, snapshots of the volume containing the new LUN cannot be taken until the LUN move operation reaches 'Moving Data' status.

If the destination is promoted late, the new LUN will be visible only after it has been fully framed. However, the LUN move job will not block the creation of snapshots of the volume containing the new LUN.

If this parameter is not specified, the destination LUN will be promoted early.

`[-max-throughput {<integer>[KB|MB|GB|TB|PB]}]` - Maximum Transfer Rate (per sec)

Optionally specifies the maximum amount of data, in bytes, that can be transferred per second in support of this operation. This mechanism can be used to throttle a transfer, to reduce its impact on the performance of the source and destination nodes.

If this parameter is not specified, throttling is not applied to the data transfer.

> The specified value will be rounded up to the nearest megabyte.

Examples

```
cluster1::> lun move start -vserver vs1 -destination-path /vol/vol2/lun2
-source-path /vol/vol1/lun1
```

Starts moving LUN *lun1* from volume *vol1* in Vserver *vs1* to *lun2* on volume *vol2* in Vserver *vs1* .

Related Links

- lun move-in-volume

lun persistent-reservation commands

lun persistent-reservation clear

Clear the SCSI-3 persistent reservation information for a given LUN

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

Description

Clears the persistent reservation for the specified LUN.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver.

{ `-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN. Examples of correct LUN paths: */vol/vol1/lun1* and */vol/vol1/qtree1/lun1*

| `-volume <volume name>` - Volume Name (privilege: advanced)

Specifies the volume.

`-lun <text>` - LUN Name (privilege: advanced)

Specifies the name of the LUN.

`[-qtree <qtree name>]` - Qtree Name (privilege: advanced) }

Specifies the qtree.

Examples

```
cluster1::*> lun persistent-reservation clear -vserver vs_1 -path
/vol/vol_1/lun_1
```

Clears the persistent reservation data for lun lun_1 in volume vol_1 for Vserver vs_1.

lun persistent-reservation show

Display the current reservation information for a given LUN

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

Description

Displays reservation information for a specified LUN in a Vserver. Unlike other show commands, the user must specify the LUN.

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.

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Specifies the Vserver.

{ `-path <path>` - LUN Path (privilege: advanced)

Specifies the path of the LUN. Examples of correct LUN paths: *`/vol/vol1/lun1`* and *`/vol/vol1/qtree1/lun1`*

| `-volume <volume name>` - Volume Name (privilege: advanced)

Specifies the volume.

`-lun <text>` - LUN Name (privilege: advanced)

Specifies the name of the LUN.

`[-qtree <qtree name>]` - Qtree Name (privilege: advanced) }

Specifies the qtree.

`[-scsi-revision {scsi2|scsi3}]` - SCSI Revision (privilege: advanced)

Selects the reservations that match this parameter value.

`[-entry-type {reservation|registration}]` - Reservation or Registration (privilege: advanced)

Selects the reservations that match this parameter value.

`[-protocol {fcp|iscsi}]` - Protocol (privilege: advanced)

Selects the reservations that match this parameter value.

`[-reservation-key <text>]` - Reservation Key (privilege: advanced)

Selects the reservations that match this parameter value.

[-reservation-type-code <text>] - Reservation Type (privilege: advanced)

Selects the reservations that match this parameter value. The possible values for SCSI-3 reservations are:

- write exclusive
- exclusive access
- write exclusive registrants only
- exclusive access registrants only
- write exclusive all registrants
- exclusive access all registrants

and for SCSI-2 are:

- regular
- third party

[-initiator-name <text>] - Initiator Name (privilege: advanced)

Selects the reservations that match this parameter value.

[-aptpl {true|false}] - Persist Through Power Loss (privilege: advanced)

Selects the reservations that match this parameter value. If `true` , the reservation will be preserved over a power loss. If `false` , it will not. This value is for SCSI-3 reservations only.

[-rtpid <integer>] - Relative Target Port ID (privilege: advanced)

Selects the reservations that match the specified Relative Target Port ID (RTPID). This applies to SCSI-3 reservations only.

[-target-wwpn <text>] - FCP Target WWPN (privilege: advanced)

Selects the reservations that match the specified World Wide Port Name (WWPN).

[-isid <text>] - Initiator Session ID (privilege: advanced)

Selects the reservations that match this parameter value.

[-tpgroup-tag <integer>] - TPGroup Tag (privilege: advanced)

Selects the reservations that match the specified target portal group tag. The tag identifies the tpgroup the reservation was made over.

[-third-party-initiator-name <text>] - Third Party Initiator Name (privilege: advanced)

Selects the reservations that match this parameter value (the initiator name that the reservation was made for). This is specific to third party reservation types, which is indicated by reservation-type-code.

[-zrto-status {none|primary-out-of-sync|primary-in-sync|secondary-out-of-sync|secondary-in-sync}] - ZRTO Status (privilege: advanced)

Selects the reservations that match this parameter value.

Examples

```
cluster1::*> lun persistent-reservation show -vserver vs_1
/vol/vol_1/lun_1
     Key                      Protocol Type              Initiator Name
     ------------------------ -------- -----------------
----------------------------
     APTPL: true
     a0:00:00:00:00:00:00:01 iscsi    write exclusive   iqn.1993-
08.org.debian:01:fa752b8a5a3a
     a0:00:00:00:00:00:00:01 iscsi    -                 iqn.1993-
08.org.debian:01:fa752b8a5a3a
     2 entries were displayed.
```

The example above displays the current reservations for lun_1 on Vserver vs_1.

lun portset commands

lun portset add

Add iSCSI/FCP LIFs to a portset

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

Description

This command adds existing iSCSI and FCP LIFs to a portset. To create a new portset, use the lun portset create command.

Use the network interface create command to create new LIFs.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-portset <text>` - Portset Name

Specifies the portset you want to add the LIFs to.

`-port-name <port_name>,…` - LIF or TPG Name

Specifies the LIF name you want to add to the portset.

Examples

```
cluster1::> portset add -vserver vs1 -portset ps1 -port-name lif1
```

Adds port lif1 to portset ps1 on Vserver vs1.

Related Links

- lun portset create
- network interface create

lun portset create

Creates a new portset

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

Description

This command creates a new portset for FCP and iSCSI. The portset name can include a maximum of 96 characters. You can add LIFs to the new portset. If you do not add a LIF to the portset, you create an empty portset. To add LIFs to an existing portset, use the lun portset add command.

After you create a portset, you must bind the portset to an igroup so the host knows which FC or iSCSI LIFs to access. If you do not bind an igroup to a portset, and you map a LUN to an igroup, then the initiators in the igroup can access the LUN on any LIF on the Vserver.

> You cannot bind an igroup to an empty portset because the initiators in the igroup would have no LIFs to access the LUN.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-portset <text>` - Portset Name

Specifies the name of the new portset. A portset name is a case-sensitive name that must contain one to 96 characters. Spaces are not allowed.

`[-port-name <port_name>,…]` - LIF Or TPG Name

Specifies the name of the logical interface that you want to add to the portset you want to create.

`{ [-protocol {mixed|fcp|iscsi}]` - Protocol

Specifies if the portset protocol type is *`fcp`* , *`iscsi`* , or *`mixed`* . The default is *`mixed`* .

`| [-f, -fcp <true>]` - FCP

Specifies FCP as the protocol type of the new portset.

`| [-i, -iscsi <true>]` - iSCSI }

Specifies iSCSI as the protocol type of the new portset.

Examples

```
cluster1::> portset create -vserver vs1 -portset ps1 -protocol mixed
```

Creates a portset *ps1* on Vserver *vs1* with the protocol type of *mixed* .

```
cluster1::> portset create -vserver vs1 -portset iscsips -protocol iscsi
```

Creates a portset *iscsips* on Vserver *vs1* with the protocol type of *iscsi* .

```
cluster1::> portset create -vserver vs1 -portset fcppc -protocol fcp
```

Creates a portset *fcppc* on Vserver *vs1* with the protocol type of *fcp* .

```
cluster1::> portset create -vserver vs1 -portset ps2 -protocol mixed -port
-name l11
```

Creates a portset *ps2* on Vserver *vs1* with the protocol type of *mixed* and LIF *l11* .

Related Links

- lun portset add

lun portset delete

Delete the portset

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

Description

This command deletes an existing portset. By default, you cannot delete a portset if it is bound to an initiator group. If a portset is bound to an initiator group, you can do one of the following:

- specify the `force` option to unbind the portset from the initiator group and delete the portset.
- use the lun igroup unbind command to unbind the portset from the initiator group. Then you can delete the portset.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-portset <text>` - Portset Name

Specifies the portset you want to delete.

`[-f, -force <true>]` - Force

Forcibly unbinds the portset from the initiator group.

Examples

```
cluster1::> portset delete -vserver vs1 -portset ps1
```

Deletes portset ps1 on Vserver vs1.

Related Links

- lun igroup unbind

lun portset remove

Remove iSCSI/FCP LIFs from a portset

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

Description

This command removes a LIF from a portset.

You cannot remove the last LIF in a portset if the portset is bound to an initiator group (igroup). To remove the last LIF in a portset, use the lun igroup unbind command to unbind the portset from the igroup. Then you can remove the last LIF in the portset.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-portset <text>` - Portset Name

Specifies the portset you want to remove a LIF from.

`-port-name <port_name>,…` - LIF or TPG Name

Specifies the LIF name you want to remove from the portset.

Examples

```
cluster1::> portset remove -vserver vs1 -portset ps1 -port-name lif1
```

Removes port lif1 from portset ps1 on Vserver vs1.

Related Links

- lun igroup unbind

lun portset show

Displays a list of portsets

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

Description

This command displays the LIFs in a portset. By default, the command displays all LIFs in all portsets.

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.

[-vserver <Vserver Name>] - Vserver Name

Selects the portsets that match this parameter value.

[-portset <text>] - Portset Name

Selects the portsets that match this parameter value.

[-uuid <UUID>] - Portset Uuid

Selects the portsets that match this parameter value.

[-port-name <port_name>,…] - LIF Or TPG Name

Selects the portsets that match this parameter value.

[-protocol {mixed|fcp|iscsi}] - Protocol

Selects the portsets that match this parameter value.

[-port-count <integer>] - Number Of Ports

Selects the portsets that match this parameter value.

[-igroups <igroup>,…] - Bound To Igroups

Selects the portsets that match this parameter value.

Examples

```
cluster1::> lun portset show
Vserver   Portset      Protocol Port Names              Igroups
--------- ------------ -------- ----------------------- ------------
vs1       ps0          mixed    lif1, lif2              igroup1
          ps1          iscsi    lif3                    igroup2
          ps2          fcp      lif4                    -
3 entries were displayed.
```

The example above displays all portsets.

```
cluster1::> lun portset show -port-count 0
Vserver   Portset      Protocol Port Names              Igroups
--------- ------------ -------- ----------------------- ------------
vs1       p1           iscsi    -                       -
```

The example above displays the portsets that contain zero LIFs.

```
cluster1::> lun portset show -protocol iscsi
Vserver   Portset      Protocol Port Names              Igroups
--------- ------------ -------- ----------------------- ------------
vs1       p1           iscsi    -                       -
vs1       iscsips      iscsi    lif1                    igroup1
2 entries were displayed.
```

The example above displays the portsets that have the iSCSI protocol.

```
cluster1::> lun portset show -port-name lif1
Vserver   Portset      Protocol Port Names              Igroups
--------- ------------ -------- ----------------------- ------------
vs1       iscsips      iscsi    lif1                    igroup1
```

The example above displays the portsets that contain LIF lif1.

lun transition commands

lun transition show

Display the status of LUN transition processing

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

Description

The `lun transition show` command displays information about the LUN transition processing status of volumes. If no parameters are specified, the command displays the following information about all volumes:

- Vserver name
- Volume name
- Transition status

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.

[-vserver <Vserver Name>] - Vserver Name (privilege: advanced)

Selects the volumes in the specified Vserver.

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

Selects the volumes with the specified name.

[-status {none|complete|failed|active}] - Transition Status (privilege: advanced)

Selects the volumes that match the specified transition status. The possible status values are:

- *active* - The volume is in an active SnapMirror transition relationship and not yet transitioned.
- *complete* - LUN transition has completed for the volume.
- *failed* - LUN transition has failed for the volume.
- *none* - The volume did not contain LUNs to transition from ONTAP 7-Mode.

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

Selects the volumes in the Vserver that matches the specified UUID.

[-node <nodename>] - Filer ID (privilege: advanced)

Selects the volumes that match the specified node.

Examples

The following example displays LUN transition information for all volumes in a Vserver named *vs1* :

```
cluster1::*> lun transition show -vserver vs1
Vserver              Volume             Transition Status
-------------------- ------------------ -----------------
vs1                  vol0               none
                     vol1               complete
                     vol2               failed
                     vol3               active
4 entries were displayed.
```

lun transition start

Start LUN Transition Processing

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

Description

The `lun transition start` command starts LUN transition for the specified volume. Normally, transition is started automatically when snapmirror break is issued for the volume, this command allows restarting in the event automatic transitioning was interrupted or failed.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

The name of the Vserver containing the volume. If only one data Vserver exists, you do not need to specify this parameter.

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

The name of the volume to restart LUN transition.

Examples

The following example starts LUN transition on a volume named *`volume1`* in a Vserver named *`vs1`* :

```
cluster1::*> lun transition start -vserver vs1 -volume volume1
```

Related Links

- snapmirror break

lun transition 7-mode delete

Delete an Untransitioned 7-Mode LUN

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

Description

The `lun transition 7-mode delete` command deletes an untransitioned LUN copied from a ONTAP 7-Mode system. This allows the admin to recover space from the volume for LUNs that may not be transitioned to ONTAP without disrupting LUNs that have transitioned, for example, if the LUN is an unsupported OS type.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

This specifies the name of the Vserver from which the LUN is to be deleted. If only one data Vserver exists, you do not need to specify this parameter.

`-path <path>` - LUN Path (privilege: advanced)

This specifies the path to the LUN to delete.

Examples

The following example deletes the LUN *`/vol/vol1/lun1`* in a Vserver named *`vs1`* :

```
cluster1::*> lun transition 7-mode delete -vserver vs1 -path
/vol/vol1/lun1
```

lun transition 7-mode show

Display the 7-Mode LUN Inventory

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

Description

The `lun transition 7-mode show` command displays information about LUNs copied from a ONTAP 7-Mode system. If no parameters are specified, the command displays the following information about all 7-Mode LUNs:

- Vserver name
- LUN path
- Operating system type
- Size
- Whether or not the LUN has been transitioned to ONTAP

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.

`[-vserver <Vserver Name>]` - Vserver Name (privilege: advanced)

Selects the 7-Mode LUNs in the specified Vserver.

`[-path <path>]` - LUN Path (privilege: advanced)

Selects the 7-Mode LUNs with the specified path.

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

Selects the 7-Mode LUNs that match the specified volume.

`[-ostype <LUN Operating System Format>]` - OS Type (privilege: advanced)

Selects the 7-Mode LUNs that match the specified operating system type.

`[-size <size>]` - LUN Size (privilege: advanced)

Selects the 7-Mode LUNs that match the specified size.

`[-prefix-size <size>]` - Prefix Stream Size (privilege: advanced)

Selects the 7-Mode LUNs that match the specified prefix stream size.

[-suffix-size <size>] - Suffix Stream Size (privilege: advanced)

Selects the 7-Mode LUNs that match the specified suffix stream size.

[-serial <text>] - Serial Number (privilege: advanced)

Selects the 7-Mode LUNs that match the specified serial number for ONTAP. LUNs where `is-transitioned` is *`false`* do not have a serial number assigned for ONTAP.

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

Selects the 7-Mode LUNs that match the specified UUID for ONTAP. LUNs where `is-transitioned` is *`false`* do not have a UUID assigned for ONTAP.

[-serial-7-mode <text>] - 7-mode Serial Number (privilege: advanced)

Selects the 7-Mode LUNs that match the specified serial number from 7-Mode.

[-is-transitioned {true|false}] - Transition Complete (privilege: advanced)

Selects the 7-Mode LUNs that match the specified transition state. LUNs where this value is *`true`* have been transitioned and are available to be mapped for client access. LUNs where this value is *`false`* have not yet been transitioned and may not be mapped.

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

Selects the 7-Mode LUNs that match the specified Vserver UUID.

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

Selects the 7-Mode LUNs that match the specified node name.

Examples

The following example displays a summary of all 7-Mode LUNs for the volume *`vol1`* in a Vserver named *`vs1`* :

```
cluster1::*> lun transition 7-mode show -vserver vs1 -volume vol1
Vserver   Path                          Type     Size   Transitioned
--------- ---------------------------- -------- ------ ------------
vs1       /vol/vol1/lun1                linux      10MB false
          /vol/vol1/lun2                linux     500MB true
          /vol/vol1/lun3                linux     500MB true
8 entries were displayed.
```

The following example displays detailed information for the 7-Mode LUN *`/vol/vol1/lun2`* in a Vserver named *`vs1`* :

```
cluster1::*> lun transition 7-mode show -vserver vs1 -path /vol/vol1/lun2
Vserver Name: vs1
            LUN Path: /vol/vol1/lun2
         Volume Name: vol1
             OS Type: linux
            LUN Size: 500MB
  Prefix Stream Size: 0
  Suffix Stream Size: 0
       Serial Number: BCVvv$DLZu8g
                UUID: f53d603b-9663-4567-9680-95c1a9cc6d9e
7-mode Serial Number: C4eqKotPI8Ui
 Transition Complete: true
        Vserver UUID: be4cc135-163f-11e3-931f-123478563412
                Node: cluster-01
```

metrocluster commands

metrocluster configure

Configure MetroCluster and start DR mirroring for the node and its DR group

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

Description

The `metrocluster configure` command creates a MetroCluster configuration on either all the nodes in both MetroCluster clusters or solely on nodes in a DR group. The command configures a HA partner, DR partner, and a DR auxiliary partner for the nodes and also starts NVRAM mirroring between the configured DR partner nodes.

In MetroCluster, a DR group is a group of four nodes, two in each of the MetroCluster clusters:

- In the local cluster, a node and its HA partner,
- In the peer cluster, a node and its HA partner. These nodes are DR partners to the nodes in the local cluster.

In a two node MetroCluster configuration, a DR group is a group of two nodes, one in each of the MetroCluster clusters.

There can be several DR groups in the MetroCluster configuration. MetroCluster provides synchronous DR protection to all data sets belonging to nodes within a properly configured DR group.

Without the `-node` parameter, the `metrocluster configure` command configures all the DR groups in both the MetroCluster clusters.

With the `-node`*`mynode`* parameter, the command configures both the *`mynode`* node and its HA partner node from the local cluster, and its DR partner and DR auxiliary partner from the peer cluster.

Before running the `metrocluster configure` command, the aggregates and Vservers on each node must be prepared for the MetroCluster configuration. Each node should have:

- At least one non-root, mirrored aggregate of size greater than 10GB. This non-root aggregate should not have any volumes in it.
- No other non-root aggregates. Any other non-root, unmirrored aggregates and volumes should be deleted.
- No Vservers other than Vservers of type "node" or "admin." Any Vservers that are not of type "node" or "admin" should be deleted.
- A mirrored and healthy root aggregate.

After the command is successful all nodes in the local and remote clusters will have HA, DR, and DR auxiliary partners and NVRAM mirroring between the DR partners will be turned on. The same conditions apply before running the `metrocluster configure-node`*`mynode`*` command, except that only one DR group is configured.

For a MetroCluster over IP configuration, the `metrocluster configuration-settings` commands must be completed before using the `metrocluster configure` command. The commands required to be completed are:

- metrocluster configuration-settings dr-group create
- metrocluster configuration-settings interface create
- metrocluster configuration-settings connection connect

Parameters

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

This optional parameter specifies the name of a single node in the local cluster. The command creates MetroCluster configuration on the local node specified by this parameter and the three other nodes belonging to the same DR group.

`[-refresh {true|false}]` - Refresh Configuration (privilege: advanced)

This optional parameter specifies if the node partner configuration steps should be done again. Not specifying this parameter will cause the MetroCluster configuration to continue using the current node partner information.

`[-allow-with-one-aggregate {true|false}]` - Override the Two Data Aggregates Requirement (privilege: advanced)

This optional parameter specifies if MetroCluster configuration should be allowed with only one data aggregate in each cluster. This option has no effect if two or more aggregates are present.

Examples

The following example shows the creation of the MetroCluster configuration for a single DR group:

```
clusA::> metrocluster show
    Cluster                         Configuration State    Mode
    ------------------------------- ----------------------
------------------------
     Local: clusA                   not-configured         -
    Remote: clusB                   not-configured         -
clusA::> metrocluster node show
    DR                                Configuration  DR
    Group Cluster Node                State          Mirroring Mode
    ----- ------- ------------------- -------------- ---------
--------------------
    -     clusA   clusA-01            ready to configure
                                                     -         -
                  clusA-02            ready to configure
                                                     -         -
                  clusA-03            ready to configure
                                                     -         -
                  clusA-04            ready to configure
                                                     -         -
    4 entries were displayed.
clusA::> metrocluster configure -node clusA-01
    [Job 45] Job succeeded: Configure is successful
clusA::> metrocluster show
    Cluster                         Configuration State    Mode
    ------------------------------- ----------------------
------------------------
     Local: clusA                   partially-configured   normal
    Remote: clusB                   partially-configured   normal
clusA::> metrocluster node show
    DR                                Configuration  DR
    Group Cluster Node                State          Mirroring Mode
    ----- ------- ------------------- -------------- ---------
--------------------
    -     clusA   clusA-03            ready to configure
                                                     -         -
                  clusA-04            ready to configure
                                                     -         -
    1     clusA   clusA-01            configured     enabled   normal
                  clusA-02            configured     enabled   normal
          clusB   clusB-01            configured     enabled   normal
                  clusB-02            configured     enabled   normal
    6 entries were displayed.
```

The following example shows the creation of the MetroCluster configuration for all DR groups:

```
clusA::> metrocluster show
    Cluster                         Configuration State    Mode
    ------------------------------- ----------------------
------------------------
     Local: clusA                   not-configured         -
    Remote: clusB                   not-configured         -
clusA::> metrocluster node show
    DR                                Configuration  DR
    Group Cluster Node                State          Mirroring Mode
    ----- ------- ------------------- -------------- ---------
--------------------
    -     clusA   clusA-01            ready to configure
                                                     -         -
                  clusA-02            ready to configure
                                                     -         -
                  clusA-03            ready to configure
                                                     -         -
                  clusA-04            ready to configure
                                                     -         -
    4 entries were displayed.
clusA::> metrocluster configure
    [Job 45] Job succeeded: Configure is successful
clusA::> metrocluster show
    Cluster                         Configuration State    Mode
    ------------------------------- ----------------------
------------------------
     Local: clusA                   configured             normal
    Remote: clusB                   configured             normal
clusA::> metrocluster node show
    DR                                Configuration  DR
    Group Cluster Node                State          Mirroring Mode
    ----- ------- ------------------- -------------- ---------
--------------------
    1     clusA   clusA-01            configured     enabled   normal
                  clusA-02            configured     enabled   normal
          clusB   clusB-01            configured     enabled   normal
                  clusB-02            configured     enabled   normal
    2     clusA   clusA-03            configured     enabled   normal
                  clusA-04            configured     enabled   normal
          clusB   clusB-03            configured     enabled   normal
                  clusB-04            configured     enabled   normal
    8 entries were displayed.
```

Related Links

- metrocluster configuration-settings dr-group create
- metrocluster configuration-settings interface create
- metrocluster configuration-settings connection connect

metrocluster heal

Heal DR data aggregates and DR root aggregates

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

Description

The `metrocluster heal` command heals DR data aggregates and DR root aggregates in preparation for a DR switchback. You must issue this command twice to complete the two phases of the healing process: first to heal the aggregates by resynchronizing the mirrored plexes and then to heal the root aggregates by switching them back to the disaster site. The DR partner nodes must be powered off and remote disk shelves must be powered on before running this command.

Parameters

`-phase {aggregates|root-aggregates}` - MetroCluster Healing Phase

This parameter specifies the healing phase. The first phase, *`aggregates`* , heals aggregates by resynchronizing mirrored plexes. The second phase, *`root-aggregates`* , heals the root aggregates of partner nodes. Healing root aggregates switches them back to the disaster site, allowing the site to boot.

`[-override-vetoes <true>]` - Override All Soft Vetoes

This optional parameter overrides almost all heal operation soft vetoes. If this optional parameter is set to true, the system overrides subsystem soft vetoes that might prevent the heal operation. Hard vetoes cannot be overridden and can still prevent the switchback operation.

Examples

The following example performs the healing of both the aggregates and root aggregates:

```
cluster1::> metrocluster heal -phase aggregates
        [Job 136] Job succeeded: Heal Aggregates is successful
cluster1::> metrocluster heal -phase root-aggregates
        [Job 137] Job succeeded: Heal Root Aggregates is successful
```

metrocluster modify

Modify MetroCluster configuration options

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

Description

The `metrocluster modify` command modifies MetroCluster parameters for nodes in the MetroCluster configuration.

Parameters

{ -auto-switchover-failure-domain <MetroCluster AUSO Failure Domain> - Cluster Level AUSO Option

This parameter specifies the configuration of automatic switchover.

The parameter values are:

- auso-on-cluster-disaster - triggers an unplanned switchover if all nodes in a DR cluster are down.
- auso-on-dr-group-disaster - triggers an unplanned switchover if both nodes of a DR group are down.
- auso-disabled - automatic switchover is disabled.

A change to the parameter affects only the local cluster where the `metrocluster modify` command is used.

| -enable-unmirrored-aggr-deployment {true|false} - Enable or Disable Unmirrored Aggregate Deployment

Enable or disable unmirrored aggregates deployment

| -is-encryption-enabled {true|false} - Is Encryption Enabled

This parameter is used to enable or disable NVLog and Storage data encryption in MetroCluster over IP configuration.

-dr-group-id <integer> - DR Group ID

This parameter identifies the DR group where encryption is being enabled or disabled.

| [-allow-auto-forced-switchover {true|false}] - Allow Automatic Forced Switchover (privilege: advanced)

This parameter is used to enable automatic forced switchover on node failures. The *`allow-auto-forced-switchover`* parameter is only supported on a MetroCluster over IP configuration.

All nodes in a MetroCluster configuration must have this option enabled to enable automatic forced switchover on failure.

| -node-name {<nodename>|local} - Node to Change the Option On }

This parameter is used to specify the node in the cluster for which the parameter needs to be modified.

[-automatic-switchover-onfailure <true>] - Node Level AUSO Option (privilege: advanced) }

This parameter is used to enable automatic switchover on node failures. The *`automatic-switchover-onfailure`* parameter is not supported on a MetroCluster over IP configuration.

All nodes in a MetroCluster configuration must have this option enabled to enable automatic switchover on failure.

Examples

The following example shows the output of MetroCluster modification done on a node:

```
clusA::*> metrocluster modify -node-name clusA-01 -node-object-limit on
  [Job 168] Job succeeded: Modify is successful
clusA::*> metrocluster modify -node-name clusA-01 -automatic-switchover
-onfailure false
  [Job 308] Job succeeded: Modify is successful
clusA::> metrocluster modify -auto-switchover-failure-domain auso-on-
cluster-disaster
  [Job 308] Job succeeded: Modify is successful
clusA::> metrocluster modify -is-encryption-enabled true -dr-group-id 1
  [Job 206] Job succeeded: Modify is successful
```

metrocluster remove-dr-group

Remove a DR group from a MetroCluster configuration with multiple DR groups

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

Description

The `metrocluster remove-dr-group` command removes a DR group from a MetroCluster configuration with multiple DR groups.

The command removes the local nodes in the identified DR group from the MetroCluster configuration. The command must be repeated on the partner cluster to remove the remote nodes in the identified DR group.

A DR group cannot be removed when the MetroCluster configuration has only a single DR group. Two or more DR groups must be configured. The command does not allow a user from unconfiguring the entire MetroCluster. Contact technical support to unconfigure the entire MetroCluster.

The following preparation steps must be completed on the local and partner clusters before removing a DR group.

- Move all data volumes to another DR group.
- Move all MDV_CRS metadata volumes to another DR group.
- Delete all MDV_aud metadata volumes that may exist in the DR group to be removed.
- Delete all data aggregates in the DR group to be removed. Root aggregates are not deleted.
- Migrate all data LIFs to home nodes in another DR group.
- Migrate the cluster management LIF to a home node in another DR group. Node management and inter-cluster LIFs are not migrated.
- Transfer epsilon to a node in another DR group.

The command is vetoed if the preparation steps are not completed on the local and partner clusters.

Parameters

`-dr-group-id <integer>` - DR Group ID (privilege: advanced)

This parameter identifies the DR group to be removed.

Examples

```
cluster_A::*> metrocluster remove-dr-group -dr-group-id 1

Warning: Nodes in the DR group that are removed from the MetroCluster
         configuration will lose their disaster recovery protection.
Local nodes "node_A1, node_A2" will be removed from the Metro-
         Cluster configuration. You must repeat the operation on the
         partner cluster "cluster_B" to remove the remote nodes in the
         DR group.
Do you want to continue? {y|n}: y

Info: The following preparation steps must be completed on the local and
      partner clusters before removing a DR group.
1. Move all data volumes to another DR group.
      2. Move all MDV_CRS metadata volumes to another DR group.
      3. Delete all MDV_aud metadata volumes that may exist in the DR
      group to be removed.
      4. Delete all data aggregates in the DR group to be removed. Root
      aggregates are not deleted.
      5. Migrate all data LIFs to home nodes in another DR group.
      6. Migrate the cluster management LIF to a home node in another DR
      group.  Node management and inter-cluster LIFs are not migrated.
      7. Transfer epsilon to a node in another DR group.
The command is vetoed if the preparation steps are not completed
      on the local and partner clusters.
Do you want to continue? {y|n}: y
[Job 80] Job succeeded: Remove DR Group is successful.
```

metrocluster show

Display MetroCluster configuration information

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

Description

The `metrocluster show` command displays configuration information for the pair of clusters configured in MetroCluster.

This command displays the following details about the local cluster and the DR partner cluster:

- Configuration State: This field specifies the configuration state of the cluster.
- Mode: This field specifies the operational mode of the cluster.
- AUSO Failure Domain: This field specifies the AUSO failure domain of the cluster.

Parameters

[-periodic-check-status]

If this option is used the MetroCluster periodic check status is displayed.

Examples

The following example shows the output of the command before MetroCluster configuration is done:

```
clusA::> metrocluster show
  Cluster                          Entry Name           State
  ------------------------------- ------------------
---------------------
  Local: clusA
          Configuration State  not-configured
                   Mode  -
          AUSO Failure Domain  -
  Remote: clusB
          Configuration State  not-configured
                   Mode  -
          AUSO Failure Domain  -
```

The following example shows the output of the command after MetroCluster configuration is done only for some DR groups:

```
clusA::> metrocluster show
  Cluster                          Entry Name           State
  ------------------------------- -------------------
---------------------
  Local: clusA
          Configuration State  partially-configured
                   Mode  -
          AUSO Failure Domain  -
  Remote: clusB
          Configuration State  partially-configured
                   Mode  -
          AUSO Failure Domain  -
```

The following example shows the output of the command after MetroCluster configuration is done:

```
clusA::> metrocluster show
  Cluster                          Entry Name           State
  ------------------------------- -------------------
---------------------
  Local: clusA
          Configuration State  configured
                   Mode  normal
          AUSO Failure Domain  auso-on-cluster-disaster
  Remote: clusB
          Configuration State  configured
                   Mode  normal
          AUSO Failure Domain  auso-on-cluster-disaster
```

The following example shows the output of the command in switchover mode:

```
clusA::> metrocluster show
  Cluster                          Entry Name           State
  ------------------------------- -------------------
---------------------
  Local: clusA
          Configuration State  configured
                   Mode  switchover
          AUSO Failure Domain  auso-on-cluster-disaster
  Remote: clusB
          Configuration State  not-reachable
                   Mode  -
          AUSO Failure Domain  not-reachable
```

The following example shows the output of the command when -peridiodic-check-status option is used:

```
clusA::> metrocluster show -periodic-check-status
       Cluster                         Periodic Check Enabled
       ------------------------------ ----------------------
        Local: clusA                   true
       Remote: clusB                   true
```

metrocluster switchback

Switch back storage and client access

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

Description

```
The `metrocluster switchback` command initiates the switchback of storage
and client access from nodes in the DR site to their home nodes. The home
nodes and storage shelves must be powered on and reachable by nodes in the
DR site. The xref:{relative_path}metrocluster-heal.html[metrocluster
heal]`-phase```_aggregates_`` and xref:{relative_path}metrocluster-
heal.html[metrocluster heal]`-phase```_root-aggregates_`` commands must
have successfully completed before running the `metrocluster switchback`
command.
```

Parameters

`[-f, -override-vetoes <true>]` - Override All Soft Vetoes

This optional parameter overrides all switchback operation soft vetoes. If this optional parameter is used, the system overrides subsystem soft vetoes that might prevent the switchback operation. Hard vetoes cannot be overridden and can still prevent the switchover operation.

`[-simulate <true>]` - Simulate Switchback (privilege: advanced)

If this optional parameter is used, the system runs a simulation of the switchback operation to make sure all the prerequisites for the operation are met. This parameter cannot be used with switchback operations performed for switching back left-behind aggregates or for retrying a partially successful switchback.

Examples

The following is an example of how to start the switchback operation.

```
clusA::> metrocluster switchback
```

Related Links

- metrocluster heal

metrocluster switchover

Switch over storage and client access

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

Description

```
The `metrocluster switchover` command initiates the switchover of storage
and client access from the source cluster to the disaster recovery (DR)
site. This command is to be used after a disaster that renders all the
nodes in the source cluster unreachable and powered off. It can also be
used for negotiated switchover when the outage of the source cluster is
anticipated as in cases such as disaster recovery testing or a site going
offline for maintenance. If a switchover operation previously failed on
certain nodes on the DR site then issuing the command retries the
operation on all of those nodes.
```

Parameters

{ [-simulate <true>] - Simulate Negotiated Switchover (privilege: advanced)

If this optional parameter is used, the system runs a simulation of the negotiated switchover operation to make sure all the prerequisites for the operation are met. This parameter cannot be used with switchover with the `-forced-on-disaster` parameter.

{ [-controller-replacement <true>] - Indicate Controller Replacement

If this optional parameter is used, the nodes in the DR site will be halted during the negotiated switchover operation, to facilitate controller replacement.

| [-forced-on-disaster <true>] - Force Switchover on Disaster

This optional parameter forces a switchover on disaster. This parameter should be used if all the nodes on the disaster stricken site are powered off and unreachable. In the absence of this parameter, the command attempts to perform a negotiated switchover operation.

[-force-nvfail-all <true>] - Sets in-nvfailed-state on All Volumes (privilege: advanced)

If this parameter is used, the switchover command will set the in-nvfailed-state parameter to true for all volumes being switched over and will set the `-dr-force-nvfail` parameter to true for any volumes that do not already have it enabled. This parameter has no effect when performing a negotiated switchover.

[-retry-failed-nodes <Node name>,…] - Nodes to Switchover }

This optional parameter takes the list of nodes that previously failed the switchover operation and it retries the switchover operation on each of the nodes. This parameter is applicable only for a switchover with the `-forced-on-disaster` parameter.

[-override-vetoes <true>] - Override All Soft Vetoes

This optional parameter overrides all switchover operation soft vetoes. If this parameter is used, the system overrides all subsystem soft vetoes that might prevent the switchover operation. Hard vetoes cannot be overridden and can still prevent the switchover operation.

Examples

When a disaster strikes one site, the `metrocluster switchover` command is issued on the disaster recovery site as follows:

```
cluster1::> metrocluster switchover -forced-on-disaster true

Warning: MetroCluster switchover is a Disaster Recovery operation that
could
         cause some data loss. The cluster on the other site must either
be
         prevented from serving data or be simply powered off (nodes and
disk
         shelves)
         The following nodes ( cluster1-01 cluster1-02 ) will participate
in
         the switchover operation
Do you want to continue? {y|n}: y
Queued job. Use 'metrocluster operation show' to check status of the DR
operation.
cluster1::> metrocluster operation show
  Operation: switchover
      State: successful
 Start time: 10/3/2013 22:11:47
   End time: 10/3/2013 22:11:53
     Errors: -
```

metrocluster check commands

metrocluster check disable-periodic-check

Disable Periodic Check

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

Description

The `metrocluster check disable-periodic-check` command disables the periodic checking of the MetroCluster configuration.

After this command is run, the MetroCluster Check job will be prevented from periodically checking the configuration for errors.

Parameters

Examples

```
clusA::> metrocluster check disable-periodic-check
```

metrocluster check enable-periodic-check

Enable Periodic Check

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

Description

The `metrocluster check enable-periodic-check` command enables the periodic checking of the MetroCluster configuration.

After this command is run, the MetroCluster Check job will able to run in the background and periodically check the configuration for errors.

Parameters

Examples

```
clusA::> metrocluster check enable-periodic-check
```

metrocluster check run

Check the MetroCluster setup

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

Description

The `metrocluster check run` command performs checks on the MetroCluster configuration and reports configuration errors if any.

To run this command, at least one DR group needs to be configured. The command checks the following parts of the configuration:

Node Configuration:

- node-reachable: This check verifies that the node is reachable.
- metrocluster-ready: This check verifies that the node is ready for MetroCluster configuration.
- local-ha-partner: This check verifies that the HA partner node is in the same cluster.
- ha-mirroring-on: This check verifies that HA mirroring for the node is configured.
- symmetric-ha-relationship: This check verifies that the relationship between the node and its HA partner is symmetric.
- remote-dr-partner: This check verifies that the DR partner node is in the remote cluster.
- dr-mirroring-on: This check verifies that DR mirroring for the node is configured.
- symmetric-dr-relationship: This check verifies that the relationship between the node and its DR partner is symmetric.
- remote-dr-auxiliary-partner: This check verifies that the DR auxiliary partner node is in the remote cluster.

- symmetric-dr-auxiliary-relationship: This check verifies that the relationship between the node and its DR auxiliary partner is symmetric.
- storage-failover-enabled: This check verifies that storage failover is enabled.
- has-intercluster-lif: This check verifies that the node has an intercluster LIF.
- node-object-limit: This check verifies that the node object limit option for the node is turned on.

Aggregate Configuration:

- mirroring-status: This check verifies that the aggregate is mirrored.
- disk-pool-allocation: This check verifies that the disks belonging to this aggregate have been correctly allocated to the right pools.

At the end of the check the command displays a summary of the results. This summary output can be viewed again by running metrocluster check show . If any of the rows in this output show any warnings more details can be viewed by running the metrocluster check show command for that component.

Parameters

`[-skip-dr-simulation {true|false}]` - Skip the DR Readiness Checks (privilege: advanced)

If this optional parameter is set to true, the switchover and switchback simulations are not run.

Examples

The following example shows the execution of the command when there are no warnings:

```
clusA::> metrocluster check run

        Last Checked On: 4/9/2014 20:11:46
Component           Result
        ------------------- ---------
        nodes               ok
        clusters            ok
        lifs                ok
        config-replication  ok
        aggregates          ok
5 entries were displayed.
Command completed. Use the "metrocluster check show -instance" command or
sub-commands in "metrocluster check" directory for detailed results.
```

The following example shows the execution of the command when there are some warnings:

```
clusA::> metrocluster check run
Last Checked On: 4/9/2014 20:11:46
Component           Result
        ------------------- ---------
        nodes               warning
        clusters            ok
        lifs                ok
        config-replication  ok
        aggregates          ok
5 entries were displayed.
Command completed. Use the "metrocluster check show -instance" command or
sub-commands in "metrocluster check" directory for detailed results.
```

Related Links

- metrocluster check show

metrocluster check show

Show the results of the last instance of MetroCluster check

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

Description

The `metrocluster check show` command displays the results of the metrocluster check run command.

This command displays the high-level verification results for each of the components. If there are any errors for a component, running the show command for that component (for example metrocluster check node show or metrocluster check aggregate show) will display more information about the warning.

> Note that this command does not run the checks but only displays the results of checks. To look at the latest results, run the metrocluster check run command and then run this command.

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.

[-timestamp <MM/DD/YYYY HH:MM:SS>] - Time of Check

This is the time at which the metrocluster check run command was last run in this cluster and these results were produced. If this parameter is specified, only rows with this timestamp will be displayed.

[-component <MetroCluster Check Components>] - Name of the Component

This is the name of the component. If this parameter is specified, only rows with this component will be displayed.

[-result {ok|warning|not-run|not-applicable}] - Result of the Check

This is the result of the check for the component. If this parameter is specified, only rows with this result will be displayed.

[-additional-info <text>] - Additional Information/Recovery Steps

This is the additional info for the verification for this component. This field will have detailed information about the warning and recovery steps. If this parameter is specified, only rows with this additional info will be displayed.

Examples

The following example shows the execution of the command when there are no warnings:

```
clusA::> metrocluster check show
cked On: 4/9/2014 20:11:46
t           Result
------------------- ---------
nodes               ok
clusters            ok
lifs                ok
config-replication  ok
aggregates          ok
connections         ok
s were displayed.
```

The following example shows the execution of the command when there are some warnings:

```
clusA::> metrocluster check show
cked On: 4/9/2014 20:11:46
t           Result
------------------- ---------
nodes               warning
clusters            ok
lifs                ok
config-replication  ok
aggregates          ok
connections         ok
s were displayed.
```

The following example shows the execution of the command with -instance option:

```
clusA::> metrocluster check show -instance
Time of Check: 4/9/2014 20:12:36
                        Name of the Component: nodes
                          Result of the Check: warning
        Additional Information/Recovery Steps:
Time of Check: 4/9/2014 20:12:36
                        Name of the Component: cluster
                          Result of the Check: ok
        Additional Information/Recovery Steps:
Time of Check: 4/9/2014 20:12:36
                        Name of the Component: lifs
                          Result of the Check: ok
        Additional Information/Recovery Steps:
Time of Check: 4/9/2014 20:12:36
                        Name of the Component: config-replication
                          Result of the Check: ok
        Additional Information/Recovery Steps:
Time of Check: 4/9/2014 20:12:36
                        Name of the Component: aggregates
                          Result of the Check: warning
        Additional Information/Recovery Steps:
Time of Check: 4/9/2014 20:12:36
                        Name of the Component: connections
                          Result of the Check: ok
        Additional Information/Recovery Steps:
        6 entries were displayed.
```

Related Links

- metrocluster check run
- metrocluster check node show
- metrocluster check aggregate show

metrocluster check aggregate show

Show results of MetroCluster check for aggregates

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

Description

The `metrocluster check aggregate show` command displays the results of aggregate checks performed by the metrocluster check run command.

The command verifies the following aspects of the configuration of all aggregates in MetroCluster:

- mirroring-status: This check verifies that the aggregate is mirrored.
- disk-pool-allocation: This check verifies that the disks belonging to this aggregate have been correctly allocated to the right pools.

Additional information about the warnings (if any) and recovery steps can be viewed by running the command with the -instance option.

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 <Node name>] - Node Name

This is the name of the node for which the check was run. If this parameter is specified, only rows with this node will be displayed.

[-aggregate <aggregate name>] - Name of the Aggregate

This is the name of the aggregate for which the check was run. If this parameter is specified, only rows with this aggregate will be displayed.

[-check <MetroCluster Aggregate Check>] - Type of Check

This is the type of the check performed. If this parameter is specified, only rows with this check will be displayed.

[-cluster <Cluster name>] - Name of Cluster

This is the name of the cluster the node belongs to. If this parameter is specified, only rows with this cluster will be displayed.

[-result {ok|warning|not-run|not-applicable}] - Result of the Check

This is the result of the check. If this parameter is specified, only rows with this result will be displayed.

[-additional-info <text>,…] - Additional Information/Recovery Steps

This is additional information about the check. This field has more information and recovery steps for the warning. If this parameter is specified, only rows with this additional info will be displayed.

Examples

The following example shows the execution of the command in a MetroCluster configuration with two nodes per cluster:

```
clusA::> metrocluster check aggregate show

Last Checked On: 4/9/2014 20:11:46

Node                  Aggregate             Check                 Result
--------------------- --------------------- ---------------------
---------
clusA-01              a1_required_data_aggr
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
                      aggr0_a1
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
clusA-02              a2_required_data_aggr
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
                      aggr0_a2
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
clusB-01              b1_required_data_aggr
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
                      aggr0_b1
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
clusB-02              aggr0_b2
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
                      b2_required_data_aggr
                                            mirroring-status      ok
                                            disk-pool-allocation  ok
16 entries were displayed.
```

The following example shows the execution of the command with -instance option:

```
clusA::> metrocluster check aggregate show -instance
Node Name: clusA-01
                Name of the Aggregate: a1_required_data_aggr_1
                        Type of Check: mirroring-status
                      Name of Cluster: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusA-01
                Name of the Aggregate: a1_required_data_aggr_1
                        Type of Check: disk-pool-allocation
```

```
                      Name of Cluster: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusA-01
                Name of the Aggregate: a1_required_data_aggr_2
                        Type of Check: mirroring-status
                      Name of Cluster: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusA-01
                Name of the Aggregate: a1_required_data_aggr_2
                        Type of Check: disk-pool-allocation
                      Name of Cluster: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusA-01
                Name of the Aggregate: aggr0_a1
                        Type of Check: mirroring-status
                      Name of Cluster: clusA
                  Result of the Check: warning
Additional Information/Recovery Steps: Root aggregate "aggr0_a1" is un-
mirrored. Root aggregates should be mirrored in a MetroCluster
configuration.
Node Name: clusA-01
                Name of the Aggregate: aggr0_a1
                        Type of Check: disk-pool-allocation
                      Name of Cluster: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusB-01
                Name of the Aggregate: aggr0_b1
                        Type of Check: mirroring-status
                      Name of Cluster: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusB-01
                Name of the Aggregate: aggr0_b1
                        Type of Check: disk-pool-allocation
                      Name of Cluster: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusB-01
                Name of the Aggregate: b1_required_data_aggr_1
                        Type of Check: mirroring-status
                      Name of Cluster: clusB
                  Result of the Check: ok
```

```
Additional Information/Recovery Steps: -
Node Name: clusB-01
                Name of the Aggregate: b1_required_data_aggr_1
                        Type of Check: disk-pool-allocation
                      Name of Cluster: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusB-01
                Name of the Aggregate: b1_required_data_aggr_2
                        Type of Check: mirroring-status
                      Name of Cluster: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps: -
Node Name: clusB-01
                Name of the Aggregate: b1_required_data_aggr_2
                        Type of Check: disk-pool-allocation
                      Name of Cluster: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps: -
12 entries were displayed.
```

Related Links

- metrocluster check run

metrocluster check cluster show

Show results of MetroCluster check for the cluster components

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

Description

The `metrocluster check cluster show` command displays the results of cluster checks performed by the metrocluster check run command.

The command displays the results of the following cluster configuration checks:

- negotiated-switchover-ready: This check verifies that the cluster is ready for a negotiated switchover operation.
- switchback-ready: This check verifies that the cluster is ready for a switchback operation.
- job-schedules: This check verifies that the job schedules between the local and remote clusters are consistent.
- licenses: This check verifies that the licenses between the local and remote clusters are consistent.
- periodic-check-enabled: This check verifies that the periodic MetroCluster Check Job is enabled.
- onboard-key-management: This check verifies that the Onboard Key Management hierarchies are consistent.

- external-key-management: This check verifies that the External Key Management configurations are consistent.

Additional information about the warnings (if any) and recovery steps can be viewed by running the command with the `-instance` parameter.

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.

[-check {negotiated-switchover-ready|switchback-ready|job-schedules|licenses|periodic-check-enabled|onboard-key-management|external-key-management}] - Type of Check

This is the type of the check performed. If this parameter is specified, only rows with this check will be displayed.

[-cluster <Cluster name>] - Cluster Name

This is the name of the cluster the check results apply to. If this parameter is specified, only rows matching the specified cluster will be displayed.

[-result {ok|warning|not-run|not-applicable}] - Result of the Check

This is the result of the check. If this parameter is specified, only rows with this result will be displayed.

[-additional-info <text>] - Additional Information/Recovery Steps

This is additional information about the check. This field has more information and recovery steps for the warning. If this parameter is specified, only rows with this additional info will be displayed.

Examples

The following example shows the execution of the command in a MetroCluster configuration:

```
clusA::> metrocluster check cluster show

Last Checked On: 11/29/2018 17:15:00

Cluster                 Check                                   Result
--------------------- ------------------------------------- ---------
clusA
                        negotiated-switchover-ready           not-applicable
                        switchback-ready                      not-applicable
                        job-schedules                         ok
                        licenses                              ok
                        periodic-check-enabled                ok
                        onboard-key-management                ok
                        external-key-management               ok
clusB
                        negotiated-switchover-ready           not-applicable
                        switchback-ready                      not-applicable
                        job-schedules                         ok
                        licenses                              ok
                        periodic-check-enabled                ok
                        onboard-key-management                ok
                        external-key-management               ok
14 entries were displayed.
```

The following example shows the execution of the command with the `-instance` parameter:

```
clusA::> metrocluster check cluster show -instance
Type of Check: negotiated-switchover-ready
                         Cluster Name: clusA
                  Result of the Check: not-applicable
Additional Information/Recovery Steps: Disaster recovery readiness checks
are not performed as part of periodic metrocluster check. To run these
checks, use the "metrocluster check run" command.
Type of Check: switchback-ready
                         Cluster Name: clusA
                  Result of the Check: not-applicable
Additional Information/Recovery Steps: Disaster recovery readiness checks
are not performed as part of periodic metrocluster check. To run these
checks, use the "metrocluster check run" command.
Type of Check: job-schedules
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: licenses
```

```
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: periodic-check-enabled
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: onboard-key-management
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: external-key-management
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: negotiated-switchover-ready
                         Cluster Name: clusB
                  Result of the Check: not-applicable
Additional Information/Recovery Steps: Disaster recovery readiness checks
are not performed as part of periodic metrocluster check. To run these
checks, use the "metrocluster check run" command.
Type of Check: switchback-ready
                         Cluster Name: clusB
                  Result of the Check: not-applicable
Additional Information/Recovery Steps: Disaster recovery readiness checks
are not performed as part of periodic metrocluster check. To run these
checks, use the "metrocluster check run" command.
Type of Check: job-schedules
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: licenses
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: periodic-check-enabled
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: onboard-key-management
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Type of Check: external-key-management
                         Cluster Name: clusB
                  Result of the Check: ok
```

```
Additional Information/Recovery Steps:
14 entries were displayed.
```

Related Links

- metrocluster check run

metrocluster check config-replication show-aggregate-eligibility

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

Description

The `metrocluster check config-replication show-aggregate-eligibility` command displays the MetroCluster configuration replication aggregate eligibility.

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.

[-aggregate <aggregate name>] - Aggregate

This is the aggregate name. If this parameter is specified, only rows with this aggregate will be displayed.

[-hosted-configuration-replication-volumes <volume name>,…] - Currently Hosted Configuration Replication Volumes

This is the list of the configuration replication volumes hosted on this aggregate. If this parameter is specified, only rows with these configuration replication volumes will be displayed.

[-is-eligible-to-host-additional-volumes {true|false}] - Eligibility to Host Another Configuration Replication Volume

This is the eligibility of the aggregate to host additional configuration replication volumes. If this parameter is specified, only rows with this eligibility will be displayed.

[-comment <text>] - Comment for Eligibility Status

This is a comment regarding the eligibility of the aggregate to host configuration replication volumes. If this parameter is specified, only rows with this comment will be displayed.

Examples

The following example shows the execution of the command in a MetroCluster configuration with thirteen aggregates in the cluster:

```
clusA::metrocluster check config-replication> show-aggregate-eligibility
                                                        Eligible to
Aggregate    Hosted Config Replication Vols             Host Addl Vols
Comments
------------ ------------------------------------------ --------------
--------
a0           -                                          false
Root Aggregate
a1           MDV_CRS_1bc7134a5ddf11e3b63f123478563412_A true           -
a2           MDV_CRS_1bc7134a5ddf11e3b63f123478563412_B true           -
a3           -                                          false
Unable to determine available space of aggregate
a4           -                                          false
Non-Local Aggregate
a5           -                                          false
Non-Home Aggregate
a6           -                                          false
Unable to determine mirror configuration
a7           -                                          false
Mirror configuration does not match requirement
a8           -                                          false
Disallowed Aggregate
a9           -                                          false
Insufficient Space - 10GB required
a10          -                                          false
Aggregate Offline
a11          -                                          false
Inconsistent Aggregate
a12          -                                          false
Aggregate Full
13 entries were displayed.
```

metrocluster check config-replication show-capture-status

Display MetroCluster capture status information

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

Description

The `metrocluster check config-replication show-capture-status` command indicates whether or not a configuration change that would prevent a negotiated switchover is currently being captured for replication.

Examples

The following example shows the execution of the command in a MetroCluster configuration when capture is not in progress:

```
cluster1::*> metrocluster check config-replication show-capture-status
Is Capture in Progress: false
```

metrocluster check config-replication show

Display MetroCluster config-replication status information

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

Description

The `metrocluster check config-replication show` command displays the results of MetroCluster configuration replication.

The command verifies the following aspects of MetroCluster configuration replication :

- Enabled: Verifies that MetroCluster configuration replication is enabled on the cluster.
- Running: Verifies that MetroCluster configuration replication is running on the cluster.
- Remote Heartbeat: Verifies that the MetroCluster configuration replication heartbeat with the remote cluster is healthy.
- Last Heartbeat Sent: Prints the timestamp of the last MetroCluster configuration replication heartbeat sent to the remote cluster.
- Last Heartbeat Received: Prints the timestamp of the last MetroCluster configuration replication heartbeat received from the remote cluster.
- Storage Status: Verifies that MetroCluster configuration replication storage is healthy.
- Storage In Use: Prints the location of MetroCluster configuration replication storage.
- Storage Remarks: Prints the underlying root cause for non healthy MetroCluster configuration storage.
- Vserver Streams: Verifies that MetroCluster configuration replication Vserver streams are healthy.
- Cluster Streams: Verifies that MetroCluster configuration replication Cluster streams are healthy.

Additional information about the warnings (if any) and recovery steps can be viewed by running the command with the -instance option.

Parameters

`[-instance ]`

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

Examples

The following example shows the output of `metrocluster check config-replication show` :

```
clusA::metrocluster check config-replication> show
                       Enabled: true
                       Running: true
               Remote Heartbeat: ok
           Last Heartbeat Sent: 12/12/2013 14:24:59
       Last Heartbeat Received: 12/12/2013 14:25:00
                Storage Status: ok
                Storage In Use: Cluster-wide Volume:
MDV_CRS_1bc7134a5ddf11e3b63f123478563412_A
               Storage Remarks: -
               Vserver Streams: ok
               Cluster Streams: ok
```

metrocluster check connection show

Display the check results of connections for nodes in a MetroCluster over IP configuration

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

Description

The `metrocluster check connection show` command displays the check results of connections for nodes in a MetroCluster over IP configuration.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| [-instance] }

If this parameter is specified, the command displays detailed information about all entries.

[-dr-group-id <integer>] - DR Group ID

If this parameter is specified, the command displays information for the matching DR group.

[-cluster-uuid <UUID>] - Cluster UUID

If this parameter is specified, the command displays information for the matching cluster specified by uuid.

[-cluster <Cluster name>] - Cluster Name

If this parameter is specified, the command displays information for the matching cluster.

[-node-uuid <UUID>] - Node UUID

If this parameter is specified, the command displays information for the matching node specified by uuid.

[-node <text>] - Node Name

If this parameter is specified, the command displays information for the matching nodes.

[-home-port {<netport>|<ifgrp>}] - Home Port

If this parameter is specified, the command displays information for the matching home-port.

[-relationship-type <Roles of MetroCluster Nodes>] - Relationship Role Type

If this parameter is specified, the command displays information for the matching relationship-type.

[-source-address <IP Address>] - Source Network Address

If this parameter is specified, the command displays information for the matching source address.

[-destination-address <IP Address>] - Destination Network Address

If this parameter is specified, the command displays information for the matching destination address.

[-partner-cluster-uuid <UUID>] - Partner Cluster UUID

If this parameter is specified, the command displays information for the matching partner-cluster-uuid.

[-partner-node-uuid <UUID>] - Partner Node UUID

If this parameter is specified, the command displays information for the matching partner-node-uuid.

[-partner-node <text>] - Partner Node Name

If this parameter is specified, the command displays information for the matching partner-node.

[-partner-type <text>] - Partner Relationship Type

If this parameter is specified, the command displays information for the matching partner-type.

[-config-state <text>] - Configuration State

If this parameter is specified, the command displays information for the matching config-state.

[-config-error-info <text>] - Configuration Error Information

If this parameter is specified, the command displays information for the matching config-error-info.

[-check-result {ok|warning|not-run|not-applicable}] - Check Connection Result

If this parameter is specified, the command displays information for the matching check-result.

[-check-ping-error-info <text>] - Check Connection Ping Error Info

If this parameter is specified, the command displays information for the matching check-ping-error-info.

[-check-mtu-size-error-info <text>] - Check Connection MTU Size Error Info

If this parameter is specified, the command displays information for the matching check-mtu-size-error-info.

[-check-storage-error-info <text>] - Check Connection Storage Error Info

If this parameter is specified, the command displays information for the matching check-storage-error-info.

Examples

The following example shows the output of the `metrocluster check connection show` command:

```
clusA::>  metrocluster check connection show
DR                     Source          Destination
```

```
Group Cluster Node    Network Address Network Address Partner Type Config
State
----- ------- ------- --------------- --------------- ------------
------------
1     cluster-A
              node-A1
                 Home Port: e0f
                      10.140.113.214  10.140.113.216  HA Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.214  10.140.113.218  DR Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.214  10.140.113.249  DR Auxiliary
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.215  10.140.113.217  HA Partner
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.215  10.140.113.248  DR Partner
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.215  10.140.113.25   DR Auxiliary
completed
              Check Result: ok
              node-A2
                 Home Port: e0f
                      10.140.113.216  10.140.113.214  HA Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.216  10.140.113.249  DR Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.216  10.140.113.218  DR Auxiliary
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.217  10.140.113.215  HA Partner
completed
```

```
              Check Result: ok
                 Home Port: e0g
                      10.140.113.217  10.140.113.25   DR Partner
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.217  10.140.113.248  DR Auxiliary
completed
              Check Result: ok
      cluster-B
              node-B1
                 Home Port: e0f
                      10.140.113.218  10.140.113.249  HA Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.218  10.140.113.214  DR Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.218  10.140.113.216  DR Auxiliary
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.248  10.140.113.25   HA Partner
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.248  10.140.113.215  DR Partner
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.248  10.140.113.217  DR Auxiliary
completed
              Check Result: ok
              node-B2
                 Home Port: e0f
                      10.140.113.249  10.140.113.218  HA Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.249  10.140.113.216  DR Partner
completed
              Check Result: ok
                 Home Port: e0f
                      10.140.113.249  10.140.113.214  DR Auxiliary
```

```
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.25   10.140.113.248  HA Partner
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.25   10.140.113.217  DR Partner
completed
              Check Result: ok
                 Home Port: e0g
                      10.140.113.25   10.140.113.215  DR Auxiliary
completed
              Check Result: ok
24 entries were displayed.
```

metrocluster check lif repair-placement

Repair LIF placement for the sync-source Vserver LIFs in the destination cluster

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

Description

The `metrocluster check lif repair-placement` command reruns LIF placement for those LIFs displayed by the metrocluster check lif show command. This command is expected to be run after the admin manually rectifies the LIF placement failures displayed in the metrocluster check lif show command output. The command is successful if the LIF placement rerun does not encounter any LIF placement failure. This is to be confirmed by subsequent running of the metrocluster check lif show .

Parameters

`-vserver <Vserver Name>` - sync-source Vserver Name

This is the name of the sync source Vserver that has LIF placement failures as reported by the metrocluster check lif show command. This input ensures that the command is run on the specified Vserver.

`[-lif <lif-name>]` - Logical Interface Name

This is the Logical Interface name that belongs to the sync source Vserver that has a LIF placement failure in the destination cluster as reported by the metrocluster check lif show command. This input ensures that the command is run on the specified LIF only.

Examples

The following example shows the execution of the command with a sync source Vserver and a LIF specified:

```
clusA::> metrocluster check lif repair-placement -vserver vs1.example.com
-lif fcplif1
Command completed. Run the "metrocluster check lif show" command for
results.

clusA::> metrocluster check lif repair-placement -vserver vs1.example.com
-lif iscsilif1
Command completed. Run the "metrocluster check lif show" command for
results.
```

The following example shows the execution of the command with only a sync-source Vserver specified:

```
clusA::> metrocluster check lif repair-placement -vserver  vs1.example.com

Command completed. Run the "metrocluster check lif show" command for
results.

clusA::>
```

Related Links

- metrocluster check lif show

metrocluster check lif show

Show results of MetroCluster check results for the data LIFs

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

Description

The `metrocluster check lif show` command displays the LIF placement failures in the MetroCluster configuration.

The command verifies the following aspects of the LIF placements of all the data LIFs in MetroCluster:

- lif-placed-on-dr-node: This check verifies that the LIF is placed on DR partner node.
- port-selection: This check verifies that the LIF is placed on correct port.

The LIF placement failures are mostly fabric/network connectivity issues that require manual intervention. Once the connectivity issues are resolved manually, the admin is expected to run metrocluster check lif repair-placement command to resolve the LIF placement issues for the sync source Vserver.

Additional information about the warnings (if any) and recovery steps can be viewed by running the command with the -instance option.

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.

[-cluster <Cluster name>] - Name of the Cluster

This is the name of the cluster the LIF belongs to. If this parameter is specified, only rows with this cluster will be displayed.

[-vserver <text>] - Name of the Vserver

This is the name of the Vserver in the MetroCluster configuration

[-lif <lif-name>] - Name of the Lif

This is the name of the LIF.

[-check <MetroCluster LIF placement Check>] - Description

This is the type of the check performed. If this parameter is specified, only rows with this check will be displayed.

[-result {ok|warning|not-run|not-applicable}] - Result of the Check

This is the result of the check performed. If this parameter is specified, only rows with this result will be displayed.

[-additional-info <text>] - Additional Information/Recovery Steps

This is additional information about the check. This field has more information and recovery steps for the warning. If this parameter is specified, only rows with this additional info will be displayed.

Examples

The following example shows the execution of the command in a MetroCluster configuration with two nodes per cluster:

```
clusA::>metrocluster check  lif show
Cluster                 Vserver     LIF         Check                  Result
------------------- ----------- -----------  ----------------------- ------
ClusA                vs1         a_data1      lif-placed-on-dr-node  ok
                                              port-selection         ok
                                 a_data1_inet6
                                              lif-placed-on-dr-node  ok
                                              port-selection         ok
ClusA                vs2-mc      b_data1      lif-placed-on-dr-node  ok
                                              port-selection
warning
                                 b_data1_inet6
                                              lif-placed-on-dr-node  ok
                                              port-selection
warning
ClusB                vs1-mc      a_data1      lif-placed-on-dr-node
warning
                                              port-selection         ok
                                 a_data1_inet6
                                              lif-placed-on-dr-node
warning
                                              port-selection         ok
ClusB                vs2         b_data1      lif-placed-on-dr-node  ok
                                              port-selection         ok
                                 b_data1_inet6
                                              lif-placed-on-dr-node  ok
                                              port-selection         ok
16 entries were displayed.
```

Related Links

- metrocluster check lif repair-placement

metrocluster check node show

Show results of MetroCluster check for nodes

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

Description

The `metrocluster check node show` command displays the results of node checks performed by the metrocluster check run command.

The command displays the results of the following node configuration checks:

- node-reachable: This check verifies that the node is reachable.

- metrocluster-ready: This check verifies that the node is ready for MetroCluster configuration.
- local-ha-partner: This check verifies that the HA partner node is in the same cluster.
- ha-mirroring-on: This check verifies that HA mirroring for the node is configured.
- ha-mirroring-op-state: This check verifies that the HA mirroring operation is online.
- symmetric-ha-relationship: This check verifies that the relationship between the node and its HA partner is symmetric.
- remote-dr-partner: This check verifies that the DR partner node is in the remote cluster.
- dr-mirroring-on: This check verifies that DR mirroring for the node is configured.
- dr-mirroring-op-state: This check verifies that the DR mirroring operation is online.
- symmetric-dr-relationship: This check verifies that the relationship between the node and its DR partner is symmetric.
- remote-dr-auxiliary-partner: This check verifies that the DR auxiliary partner node is in the remote cluster.
- symmetric-dr-auxiliary-relationship: This check verifies that the relationship between the node and its DR auxiliary partner is symmetric.
- storage-failover-enabled: This check verifies that storage failover is enabled.
- has-intercluster-lif: This check verifies that the node has an intercluster LIF.
- node-object-limit: This check verifies that the node object limit option for the node is turned on.
- automatic-uso: This check verifies that the automatic USO option for the node is enabled.

Additional information about the warnings (if any) and recovery steps can be viewed by running the command with the `-instance` parameter.

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 <Node name>] - Node Name

This is the name of the node for which the check was run. If this parameter is specified, only rows with this node will be displayed.

[-check <MetroCluster Node Check>] - Type of Check

This is the type of the check performed. If this parameter is specified, only rows with this check will be displayed.

[-cluster <Cluster name>] - Cluster Name

This is the name of the cluster the node belongs to. If this parameter is specified, only rows with this cluster will be displayed.

[-result {ok|warning|not-run|not-applicable}] - Result of the Check

This is the result of the check. If this parameter is specified, only rows with this result will be displayed.

[-additional-info <text>] - Additional Information/Recovery Steps

This is additional information about the check. This field has more information and recovery steps for the warning. If this parameter is specified, only rows with this additional info will be displayed.

Examples

The following example shows the execution of the command in a MetroCluster configuration with two nodes per cluster:

```
clusA::> metrocluster check node show

Last Checked On: 9/12/2016 13:47:00

Node                   Check                                  Result
---------------------- -------------------------------------- ---------
clusA-01
                       node-reachable                         ok
                       metrocluster-ready                     ok
                       local-ha-partner                       ok
                       ha-mirroring-on                        warning
                       ha-mirroring-op-state                  ok
                       symmetric-ha-relationship              warning
                       remote-dr-partner                      ok
                       dr-mirroring-on                        ok
                       dr-mirroring-op-state                  ok
                       symmetric-dr-relationship              ok
                       remote-dr-auxiliary-partner            ok
                       symmetric-dr-auxiliary-relationship    warning
                       storage-failover-enabled               ok
                       has-intercluster-lif                   ok
                       node-object-limit                      ok
                       automatic-uso                          ok
clusA-02
                       node-reachable                         ok
                       metrocluster-ready                     ok
                       local-ha-partner                       ok
                       ha-mirroring-on                        warning
                       ha-mirroring-op-state                  ok
                       symmetric-ha-relationship              warning
                       remote-dr-partner                      ok
                       dr-mirroring-on                        ok
                       dr-mirroring-op-state                  ok
                       symmetric-dr-relationship              ok
                       remote-dr-auxiliary-partner            ok
                       symmetric-dr-auxiliary-relationship    warning
                       storage-failover-enabled               ok
                       has-intercluster-lif                   ok
```

```
                      node-object-limit                     ok
                      automatic-uso                         ok
clusB-01
                      node-reachable                        ok
                      metrocluster-ready                    ok
                      local-ha-partner                      ok
                      ha-mirroring-on                       warning
                      ha-mirroring-op-state                 ok
                      symmetric-ha-relationship             warning
                      remote-dr-partner                     ok
                      dr-mirroring-on                       ok
                      dr-mirroring-op-state                 ok
                      symmetric-dr-relationship             ok
                      remote-dr-auxiliary-partner           ok
                      symmetric-dr-auxiliary-relationship  warning
                      storage-failover-enabled              ok
                      has-intercluster-lif                  ok
                      node-object-limit                     ok
                      automatic-uso                         ok
clusB-02
                      node-reachable                        ok
                      metrocluster-ready                    ok
                      local-ha-partner                      ok
                      ha-mirroring-on                       warning
                      ha-mirroring-op-state                 ok
                      symmetric-ha-relationship             warning
                      remote-dr-partner                     ok
                      dr-mirroring-on                       ok
                      dr-mirroring-op-state                 ok
                      symmetric-dr-relationship             ok
                      remote-dr-auxiliary-partner           ok
                      symmetric-dr-auxiliary-relationship  warning
                      storage-failover-enabled              ok
                      has-intercluster-lif                  ok
                      node-object-limit                     ok
                      automatic-uso                         ok
64 entries were displayed.
```

The following example shows the execution of the command with the `-instance` parameter:

```
clusA::> metrocluster check node show -instance
Node Name: clusA-01
                        Type of Check: node-reachable
                         Cluster Name: clusA
                   Result of the Check: ok
```

```
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: metrocluster-ready
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: local-ha-partner
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: ha-mirroring-on
                         Cluster Name: clusA
                  Result of the Check: warning
Additional Information/Recovery Steps: Node's HA mirroring is not active.
Enable it on using "storage failover" commands.
Node Name: clusA-01
                        Type of Check: ha-mirroring-op-state
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: symmetric-ha-relationship
                         Cluster Name: clusA
                  Result of the Check: warning
Additional Information/Recovery Steps: Partner not found. Check if node
"clusA-01's HA partner" is configured in MetroCluster.
Node Name: clusA-01
                        Type of Check: remote-dr-partner
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: dr-mirroring-on
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: dr-mirroring-op-state
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: symmetric-dr-relationship
                         Cluster Name: clusA
```

```
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: remote-dr-auxiliary-partner
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: symmetric-dr-auxiliary-relationship
                         Cluster Name: clusA
                  Result of the Check: warning
Additional Information/Recovery Steps: Partner not found. Check if node
"clusA-01's DR auxiliary partner" is configured in MetroCluster.
Node Name: clusA-01
                        Type of Check: storage-failover-enabled
                         Cluster Name: clusA
                  Result of the Check: warning
Additional Information/Recovery Steps: Node's storage failover is
disabled.  Enable using "storage failover" commands.
Node Name: clusA-01
                        Type of Check: has-intercluster-lif
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusA-01
                        Type of Check: node-object-limit
                         Cluster Name: clusA
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                        Type of Check: node-reachable
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                        Type of Check: metrocluster-ready
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                        Type of Check: local-ha-partner
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                        Type of Check: ha-mirroring-on
```

```
                         Cluster Name: clusB
                  Result of the Check: warning
Additional Information/Recovery Steps: Node's HA mirroring is not active.
Enable it on using "storage failover" commands.
Node Name: clusB-01
                       Type of Check: ha-mirroring-op-state
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                       Type of Check: symmetric-ha-relationship
                         Cluster Name: clusB
                  Result of the Check: warning
Additional Information/Recovery Steps: Partner not found. Check if node
"clusB-01's HA partner" is configured in MetroCluster.
Node Name: clusB-01
                       Type of Check: remote-dr-partner
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                       Type of Check: dr-mirroring-on
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                       Type of Check: dr-mirroring-op-state
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                       Type of Check: symmetric-dr-relationship
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                       Type of Check: remote-dr-auxiliary-partner
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                       Type of Check: symmetric-dr-auxiliary-relationship
                         Cluster Name: clusB
                  Result of the Check: warning
Additional Information/Recovery Steps: Partner not found. Check if node
"clusB-01's DR auxiliary partner" is configured in MetroCluster.
```

```
Node Name: clusB-01
                        Type of Check: storage-failover-enabled
                         Cluster Name: clusB
                  Result of the Check: warning
Additional Information/Recovery Steps: Node's storage failover is
disabled.  Enable using "storage failover" commands.
Node Name: clusB-01
                        Type of Check: has-intercluster-lif
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                        Type of Check: node-object-limit
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
Node Name: clusB-01
                        Type of Check: automatic-uso
                         Cluster Name: clusB
                  Result of the Check: ok
Additional Information/Recovery Steps:
32 entries were displayed.
```

Related Links

- metrocluster check run

metrocluster check volume show

Show results of the MetroCluster check for volumes

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

Description

The `metrocluster check volume show` command displays the results of volume checks performed by the metrocluster check run command.

The command displays the results of the following volume configuration checks:

- unmirrored-flexgroups: This check looks for FlexGroup volumes residing on unmirrored aggregates.
- mixed-flexgroups: This check looks for FlexGroup volumes residing on a mix of mirrored and unmirrored aggregates.

Additional information about the warnings, if any, and recovery steps can be viewed by running the command with the `-instance` parameter.

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.

[-vserver <vserver name>] - Vserver Name

This is the name of the vserver that contains the volume that the check results apply to. If this parameter is specified, only rows matching the specified cluster will be displayed.

[-volume <volume name>] - Volume Name

This is the name of the volume that the check results apply to. If this parameter is specified, only rows matching the specified volume will be displayed.

[-check <MetroCluster Volume Check>] - Type of Check

This is the type of the check performed. If this parameter is specified, only rows with this check will be displayed.

[-result {ok|warning|not-run|not-applicable}] - Result of the Check

This is the result of the check. If this parameter is specified, only rows with this result will be displayed.

[-additional-info <text>,…] - Additional Information/Recovery Steps

This is additional information about the check. This field has more information and recovery steps for the warning. If this parameter is specified, only rows with this additional info will be displayed.

Examples

The following example shows the execution of the command in a MetroCluster configuration:

```
clusA::*> metrocluster check volume show

Last Checked On: 7/25/2018 10:04:07

Vserver                Volume                    Check
Result
--------------------- ------------------------- ---------------------
-------
vs1                    unmirr
                                                 unmirrored-volumes
warning
vs2                    vs2UnMirrA
                                                 unmirrored-volumes
warning
2 entries were displayed.

clusA::*> metrocluster check volume show -instance
Vserver Name: vs1
                          Volume Name: unMirr
                        Type of Check: unmirrored-volumes
                   Result of the Check: warning
Additional Information/Recovery Steps: FlexGroup "unMirr" resides on
unmirrored aggregates. Parts of the FlexGroup may not be available after
an un-planned switchover.
Vserver Name: vs2
                          Volume Name: vs2UnMirrA
                        Type of Check: unmirrored-volumes
                   Result of the Check: warning
Additional Information/Recovery Steps: FlexGroup "vs2UnMirrA" resides on
unmirrored aggregates. Parts of the FlexGroup may not be available after
an un-planned switchover.
2 entries were displayed.
clusA::>
```

Related Links

- metrocluster check run

metrocluster config-replication commands

metrocluster config-replication cluster-storage-configuration modify

Modify MetroCluster storage configuration information

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

Description

The `metrocluster config-replication cluster-storage-configuration modify` command modifies the configuration of storage used for configuration replication.

Parameters

`[-disallowed-aggregates <aggregate name>,…]` - Disallowed Aggregates (privilege: advanced)

Use this parameter to set the list of storage aggregates that are not available to host storage for configuration replication.

Examples

The following example disallows two aggregates named *`aggr1`* and *`aggr2`* :

```
cluster1::*> metrocluster config-replication cluster-storage-configuration
modify -disallowed-aggregates aggr1,aggr2
```

metrocluster config-replication cluster-storage-configuration show

Display MetroCluster storage configuration information

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

Description

The `metrocluster config-replication cluster-storage-configuration show` command shows details of the configuration of the storage used for configuration replication.

The information displayed is the following:

- Disallowed Aggregates - The list of storage aggregates that are configured as not allowed to host storage areas.
- Auto-Repair - Displays *`true`* if the automatic repair of storage areas used by configuration replication is enabled.
- Auto-Recreate - Displays *`true`* if the automatic recreation of storage volumes used by configuration replication is enabled.
- Use Mirrored Aggregate - Displays *`true`* if storage areas for configuration replication are to be hosted on a mirrored aggregate.

Examples

The following is an example of the `metrocluster config-replication cluster-storage-configuration show` command:

```
cluster1::*> metrocluster config-replication cluster-storage-configuration
show
Disallowed Aggregates: -
           Auto-Repair: true
         Auto-Recreate: true
Use Mirrored Aggregate: true
```

metrocluster config-replication resync-status show

Display MetroCluster Configuration Resynchronization Status

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

Description

The `metrocluster config-replication resync-status show` command displays the state of the configuration synchronization operation between the two clusters in the MetroCluster configuration.

This command displays the following details about the local cluster and the peer cluster:

- Source: This is the source side whose configuration is being replicated to the destination side.
- Destination: This is the destination side where the configuration is being replicated to from the source side.
- State: This is the state of the synchronization operation.
- % Complete: This is completion percentage of the operation.

Examples

The following example shows the output of the command when synchronization is in progress:

```
clusterA::> metrocluster config-replication resync-status show
       Source                  Destination             State       %
Complete
       ----------------------- ----------------------- -----------
----------
       clusterA                clusterB                complete    -
       clusterB                clusterA                complete    -
```

The following example shows the output of the command when synchronization from clusB to clusA is in progress:

```
clusA::> metrocluster config-replication resync-status show
       Source                  Destination             State       %
Complete
       ----------------------- ----------------------- -----------
----------
       clusterA                clusterB                complete    -
       clusterB                clusterA                messaging   95
```

metrocluster configuration-settings commands

metrocluster configuration-settings show-status

Display the configuration settings status for a MetroCluster setup

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

Description

The `metrocluster configuration-settings show-status` command displays the configuration settings status for nodes in a MetroCluster setup. If a DR group has not been created, then status for nodes in the local cluster only are displayed.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| [-instance] }

If this parameter is specified, the command displays detailed information about all entries.

[-cluster-uuid <UUID>] - Cluster UUID

If this parameter is specified, the command displays detailed information about all nodes in the cluster matching the specified cluster-uuid.

[-cluster <Cluster name>] - Cluster Name

If this parameter is specified, the command displays detailed information about all the nodes in the specified cluster.

[-node <text>] - Node Name

If this parameter is specified, the command displays information for the matching nodes.

[-configuration-status <text>] - Configuration Settings Status

If this parameter is specified, the command displays detailed information about all nodes with the specified configuration status.

`[-config-error-info <text>]` - Configuration Error Information

If this parameter is specified, the command displays detailed information about all nodes with the specified configuration error information.

Examples

The following example shows the display of MetroCluster setup status:

```
Nodes do not have a valid platform-specific personality value (equivalent
to HAOSC parameter on non-Apollo platforms) for a MetroCluster setup.
clusA::> metrocluster configuration-settings show-status
Cluster                    Node               Configuration Settings
Status
-------------------------- ------------------
--------------------------------
   clusA                      A1                 not a MetroCluster setup
                              A2                 not a MetroCluster setup
2 entries were displayed.
MetroCluster setup uses FC links rather than IP
xref:{relative_path}clusA::> metrocluster configuration-settings show-
status
Cluster                    Node               Configuration Settings
Status
-------------------------- ------------------
--------------------------------
clusA                      A1                 not applicable for FC and
SAS
                           A2                 not applicable for FC and
SAS
2 entries were displayed.
Output of the command when MetroCluster setup uses IP links and before
`"metrocluster configuration-settings dr-group create"` command is run:
clusA::> metrocluster configuration-settings show-status
Cluster                    Node               Configuration Settings
Status
-------------------------- ------------------
--------------------------------
clusA                      A1                 ready for DR group create
                           A2                 ready for DR group create
2 entries were displayed.
Output of the command after `"metrocluster configuration-settings dr-group
create"` command is run:
clusA::> metrocluster configuration-settings show-status
Cluster                    Node               Configuration Settings
Status
-------------------------- ------------------
```

```
--------------------------------
clusA
                           A1                  ready for interface create
                           A2                  ready for interface create
clusB
                           B1                  ready for interface create
                           B2                  ready for interface create
4 entries were displayed.
Output of the command after `"metrocluster configuration-settings
interface create"` command is run for every node:
clusA::> metrocluster configuration-settings show-status
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
--------------------------------
clusA
                           A1                  ready for next interface
create
                           A2                  ready for connection connect
clusB
                           B1                  ready for connection connect
                           B2                  ready for connection connect
4 entries were displayed.
Output of the command after ` "metrocluster configuration-settings
connection connect" ` command is run:
usA::> metrocluster configuration-settings show-status
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
--------------------------------
clusA
                           A1                  completed
                           A2                  completed
clusB
                           B1                  completed
                           B2                  completed
4 entries were displayed.
Output of the command after ` "metrocluster configuration-settings
connection connect"` command is run and there are connection errors:
clusA::> metrocluster configuration-settings show-status
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
--------------------------------
clusA                      A1                  connection error
                           A2                  completed
```

```
clusB                       B1                  connection error
                            B2                  completed
4 entries were displayed.
```

metrocluster configuration-settings calibration measure

Measure latency and bandwith values

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

Description

The `metrocluster configuration-settings calibration measure` command measures the bandwidth and latency between local and remote nodes in a MetroCluster over IP configuration.

This command should not be run when the MetroCluster configuration is under a heavy load. Bandwidth measurements will attempt to fully saturate network connections to the remote cluster and may disrupt the quality of service.

Examples

```
The following example shows the output for the calibration measure command
in MetroCluster over IP configurations:
```

```
clusA::*> metrocluster configuration-settings calibration measure
Warning: This operation will attempt to fully saturate the network
connection to
     the remote cluster for measuring bandwidth. This may disrupt
           performance and should not be run when MetroCluster is under
heavy
     load. Do you want to continue? {y|n}: y
  Measuring latency and bandwidth between node "A1" (10.140.113.214) and
node "B1" (10.140.113.26) over port "e0f".
  Measuring latency and bandwidth between node "A1" (10.140.113.215) and
node "B1" (10.140.113.27) over port "e0g".
  Measuring latency and bandwidth between node "A2" (10.140.113.216) and
node "B2" (10.140.113.25) over port "e0f".
  Measuring latency and bandwidth between node "A2" (10.140.113.217) and
node "B2" (10.140.113.28) over port "e0g".
  Measurements complete. Use the "metrocluster configuration-settings
calibration show" command to display the results.
```

metrocluster configuration-settings calibration show

Display the calibration measurements for local nodes in MetroCluster over IP

configuration

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

Description

The `metrocluster configuration-settings calibration show` command shows the bandwidth and latency between local and remote nodes in a MetroCluster over IP configuration.

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.

[-calibration-id <integer>] - Calibration ID

This field specifies Calibration ID.

[-subnet <IP Address>] - IP Subnet

This field specifies subnet.

[-node-uuid <UUID>] - Node UUID

This field specifies UUID of the node.

[-home-port {<netport>|<ifgrp>}] - Home Port

This field specifies name of the port used to measure latency and bandwidth.

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

This field specifies name of the node.

[-collection-time <MM/DD/YYYY HH:MM:SS>] - Collection Time of Sample

This field specifies collection time of sample.

[-source-address <IP Address>] - Source Network Address

This field specifies IP address of the port on the node.

[-destination-address <IP Address>] - Destination Network Address

This field specifies IP address of the port on the partner node.

[-latency <double>] - Latency in ms

This field specifies latency in ms.

[-bandwidth <double>] - Bandwidth in Mb/s

This field specifies bandwidth in Mb/s.

[-num-packets-retransmitted <integer>] - Number of Packets Retransmitted

This field specifies the number of TCP packets retransmitted during bandwidth measurement.

[-pct-pkt-retrans <double>] - Percentage of packages retransmitted

This field specifies the percent of TCP packets retransmitted during bandwidth measurement.

Examples

```
The following example shows the output for the calibration show command in
MetroCluster over IP configurations:
```

```
clusA::*> metrocluster configuration-settings calibration show
                                    Home Latency Bandwidth
  ID Subnet         Node     Port (ms)    (Mb/s)    Collection Time
  -- -------------- ---------------- ---- ------- ---------
-------------------
  1  172.21.96.0    A1
                                    e0f  0.341   764       8/6/2019
16:20:07
                                    e0g  0.311   629       8/6/2019
16:20:22
        A2
                                    e0f  0.307   1265      8/6/2019
16:20:36
                                    e0g  0.307   1063      8/6/2019
16:20:50
  4 entries were displayed.
```

metrocluster configuration-settings connection check

Check the network connections between partner nodes

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

Description

The `metrocluster configuration-settings connection check` command checks the settings of a MetroCluster over IP configuration.

This command is used for MetroCluster configurations that are connected through IP links.

Parameters

[-v, -verbose <true>] - Enable Detailed Output (privilege: advanced)

This optional parameter enables verbose mode.

[-p, -ping <true>] - Check Ping Response Only (privilege: advanced)

Specify this parameter to perform reachability checks.

[-t, -throughput <true>] - Check Throughput Only (privilege: advanced)

Specify this parameter to perform throughput checks.

{ [-z, -tos <Hex String>] - Type of Service for Pings (privilege: advanced)

This optional parameter specifies the value for the Type of Service field to be used for ping or throughput checks. If specified, the ping and throughput checks will be limited to only the specified value and not cycle through all the pre-defined TOS values for those checks.

| [-D, -dscp <text>] - DSCP for Throughput or Pings (privilege: advanced)

This optional parameter specifies the DSCP value to be used for ping or throughput checks. The value can be specified as a number, hexadecimal string or as a symbolic name. For example, "32", "0x20" and "CS4" all refer to the same value. If specified, the pinng and throughput checks will be limited to only the specified value and not cycle through all the pre-defined DSCP values for these checks.

[-E, -ecn <integer>] - ECN for Pings (privilege: advanced) }

This optional parameter specifies the ECN value to be used for ping checks.

[-c, -count <integer>] - Ping Count (privilege: advanced)

This optional parameter specifies the number of ping requests to send for each destination. The default value is 1.

[-i, -interval <double>] - Ping Interval (privilege: advanced)

This optional parameter specifies the interval to use between pings. The default value is 1 second.

[-S, -source <IP Address>] - Source IP Address (privilege: advanced)

This optional parameter specifies the source IP address to use for ping or throughput check.

[-destination <IP Address>] - Destination IP Address (privilege: advanced)

This optional parameter specifies the destination IP address to use for ping or throughput checks.

[-role <Roles of MetroCluster Nodes>] - Partner Role (privilege: advanced)

This optional parameter specifies the partner role to use for ping and throughput checks. Only the nodes with the matching role will be used.

[-I, -port {<netport>|<ifgrp>}] - Interface Port Name (privilege: advanced)

This optional parameter specifies the port to use for ping and throughput checks.

Examples

```
The following example shows the output for the check command in
MetroCluster over IP configurations:
```

```
clusA::*> metrocluster configuration-settings connection check
[Job 68] Job succeeded: Connect is successful.

Begin connection check.
Start checking the partner cluster.
    Check partner cluster: PASS.
Start checking the configuration settings.
    Check configuration settings: PASS.
Start pinging the network endpoints from cluster "clusA".
    Ping network endpoints: PASS.
Start pinging the network endpoints from cluster "clusB".
    Ping network endpoints: PASS.
Start checking the network MTU sizes from cluster "clusA".
    Check network MTU sizes: PASS.
Start checking the network MTU sizes from cluster "clusB".
    Check network MTU sizes: PASS.
Start checking the network subnets from cluster "clusA".
    Check network subnets: PASS.
Start checking the network subnets from cluster "clusB".
    Check network subnets: PASS.
Start checking the storage daemons on cluster "clusA".
    Check storage daemons: PASS.
Start checking the storage daemons on cluster "clusB".
    Check storage daemons: PASS.
End of connection check.
```

metrocluster configuration-settings connection connect

Configure the network connections between partner nodes

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

Description

The `metrocluster configuration-settings connection connect` command configures the connections that mirror NV logs and access remote storage between partner nodes in a MetroCluster setup.

This command is used for MetroCluster setups that are connected though IP links. MetroCluster setups that are connected through FC links will configure the FC connections automatically.

The `metrocluster configuration-settings` commands are run in the following order to set up MetroCluster:

- metrocluster configuration-settings dr-group create ,
- metrocluster configuration-settings interface create ,
- `metrocluster configuration-settings connection connect` .

Before this command is run

- The DR groups must have been configured. Run the metrocluster configuration-settings dr-group show command to verify that every node is partnered in a DR group.
- The network logical interfaces must have been configured on every node. Use the metrocluster configuration-settings interface show command to verify that every node has network logical interfaces configured to mirror NV logs and access remote storage.

After this command completes successfully, every node will:

- Have NV log mirroring configured and mirroring disabled. NV log mirroring will be enabled by the metrocluster configure command.
- Have access to remote storage. Use the storage disk show -pool Pool1 command to view the remote disks that are hosted on DR partner nodes.

The DR groups and network logical interfaces that were configured by the `metrocluster configuration-settings` commands cannot be deleted after the connections have been configured. The metrocluster configuration-settings connection disconnect command must be run to remove the connections before the DR groups and network logical interfaces can be deleted.

Parameters

Examples

```
The following example shows configuration of connections in a MetroCluster
over IP setup:
```

```
clusA::> metrocluster configuration-settings connection connect
[Job 269] Job succeeded: Connect is successful.
clusA::>  metrocluster configuration-settings connection show
DR                    Source          Destination
Group Cluster Node    Network Address Network Address Partner Type Config
State
----- ------- ------- --------------- --------------- ------------
------------
1     clusA A1
                 Home Port: e0f
                      10.140.113.214  10.140.113.216  HA Partner
completed
                 Home Port: e0f
                      10.140.113.214  10.140.113.218  DR Partner
completed
                 Home Port: e0f
                      10.140.113.214  10.140.113.249  DR Auxiliary
completed
                 Home Port: e0g
                      10.140.113.215  10.140.113.217  HA Partner
```

```
completed
                 Home Port: e0g
                      10.140.113.215  10.140.113.248  DR Partner
completed
                 Home Port: e0g
                      10.140.113.215  10.140.113.25   DR Auxiliary
completed
                       A2
                 Home Port: e0f
                      10.140.113.216  10.140.113.214  HA Partner
completed
                 Home Port: e0f
                      10.140.113.216  10.140.113.249  DR Partner
completed
                 Home Port: e0f
                      10.140.113.216  10.140.113.218  DR Auxiliary
completed
                 Home Port: e0g
                      10.140.113.217  10.140.113.215  HA Partner
completed
                 Home Port: e0g
                      10.140.113.217  10.140.113.25   DR Partner
completed
                 Home Port: e0g
                      10.140.113.217  10.140.113.248  DR Auxiliary
completed
      clusB B2
                 Home Port: e0f
                      10.140.113.249  10.140.113.218  HA Partner
completed
                 Home Port: e0f
                      10.140.113.249  10.140.113.216  DR Partner
completed
                 Home Port: e0f
                      10.140.113.249  10.140.113.214  DR Auxiliary
completed
                 Home Port: e0g
                      10.140.113.25   10.140.113.248  HA Partner
completed
                 Home Port: e0g
                      10.140.113.25   10.140.113.217  DR Partner
completed
                 Home Port: e0g
                      10.140.113.25   10.140.113.215  DR Auxiliary
completed
                       B1
```

```
                 Home Port: e0f
                      10.140.113.218  10.140.113.249  HA Partner
completed
                 Home Port: e0f
                      10.140.113.218  10.140.113.214  DR Partner
completed
                 Home Port: e0f
                      10.140.113.218  10.140.113.216  DR Auxiliary
completed
                 Home Port: e0g
                      10.140.113.248  10.140.113.25   HA Partner
completed
                 Home Port: e0g
                      10.140.113.248  10.140.113.215  DR Partner
completed
                 Home Port: e0g
                      10.140.113.248  10.140.113.217  DR Auxiliary
completed
24 entries were displayed.
clusA::> metrocluster configuration-settings show-status
Cluster                    Node               Configuration Settings
Status
-------------------------- ------------------
---------------------------------
clusA
                           A1                 completed
                           A2                 completed
clusB
                           B1                 completed
                           B2                 completed
4 entries were displayed.
```

Related Links

- metrocluster configuration-settings dr-group create
- metrocluster configuration-settings interface create
- metrocluster configuration-settings dr-group show
- metrocluster configuration-settings interface show
- metrocluster configure
- storage disk show
- metrocluster configuration-settings connection disconnect

metrocluster configuration-settings connection disconnect

Tear down the network connections between partner nodes

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

Description

The `metrocluster configuration-settings connection disconnect` command removes the connections between nodes in a DR group that are used to mirror NV logs and access remote storage.

This command cannot be run if a node in the DR group has remote disks assigned to the node. The assigned ownership of remote disks can be removed by running the storage disk removeowner command.

The `metrocluster configuration-settings` commands are run in the following order to remove MetroCluster over IP configuration:

- `metrocluster configuration-settings connection disconnect` ,
- metrocluster configuration-settings interface delete ,
- metrocluster configuration-settings dr-group delete .

Parameters

`-dr-group-id <integer>` - DR Group ID

This parameter identifies the DR group to be disconnected.

Examples

The following example illustrates removal of connections in a four-node MetroCluster setup:

```
clusA::> metrocluster configuration-settings connection disconnect -dr
-group-id 1
[Job 270] Job succeeded: Disconnect is successful.

clusA::>  metrocluster configuration-settings show-status
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
---------------------------------
clusA
                           A1                  ready for connection connect
                           A2                  ready for connection connect
clusB
                           B1                  ready for connection connect
                           B2                  ready for connection connect
4 entries were displayed.
clusA::> metrocluster configuration-settings connection show
DR                     Source          Destination
Group Cluster Node     Network Address Network Address Partner Type Config
State
----- ------- ------- --------------- --------------- ------------
------------
```

```
1     clusA A1
                 Home Port: e0f
                      10.140.113.214  10.140.113.216  HA Partner
disconnected
                 Home Port: e0f
                      10.140.113.214  10.140.113.218  DR Partner
disconnected
                 Home Port: e0f
                      10.140.113.214  10.140.113.249  DR Auxiliary
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.217  HA Partner
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.248  DR Partner
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.25   DR Auxiliary
disconnected
                       A2
                 Home Port: e0f
                      10.140.113.216  10.140.113.214  HA Partner
disconnected
                 Home Port: e0f
                      10.140.113.216  10.140.113.249  DR Partner
disconnected
                 Home Port: e0f
                      10.140.113.216  10.140.113.218  DR Auxiliary
disconnected
                 Home Port: e0g
                      10.140.113.217  10.140.113.215  HA Partner
disconnected
                 Home Port: e0g
                      10.140.113.217  10.140.113.25   DR Partner
disconnected
                 Home Port: e0g
                      10.140.113.217  10.140.113.248  DR Auxiliary
disconnected
      clusB B2
                 Home Port: e0f
                      10.140.113.249  10.140.113.218  HA Partner
disconnected
                 Home Port: e0f
                      10.140.113.249  10.140.113.216  DR Partner
disconnected
                 Home Port: e0f
```

```
                      10.140.113.249  10.140.113.214  DR Auxiliary
disconnected
                 Home Port: e0g
                      10.140.113.25   10.140.113.248  HA Partner
disconnected
                 Home Port: e0g
                      10.140.113.25   10.140.113.217  DR Partner
disconnected
                 Home Port: e0g
                      10.140.113.25   10.140.113.215  DR Auxiliary
disconnected
                       B1
                 Home Port: e0f
                      10.140.113.218  10.140.113.249  HA Partner
disconnected
                 Home Port: e0f
                      10.140.113.218  10.140.113.214  DR Partner
disconnected
                 Home Port: e0f
                      10.140.113.218  10.140.113.216  DR Auxiliary
disconnected
                 Home Port: e0g
                      10.140.113.248  10.140.113.25   HA Partner
disconnected
                 Home Port: e0g
                      10.140.113.248  10.140.113.215  DR Partner
disconnected
                 Home Port: e0g
                      10.140.113.248  10.140.113.217  DR Auxiliary
disconnected
24 entries were displayed.
```

Related Links

- storage disk removeowner
- metrocluster configuration-settings interface delete
- metrocluster configuration-settings dr-group delete

metrocluster configuration-settings connection show

Display the connections between partner nodes in a MetroCluster setup

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

Description

The `metrocluster configuration-settings connection show` command displays the connection configuration information between the nodes in a MetroCluster setup.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| [-instance] }

If this parameter is specified, the command displays detailed information about all entries.

[-dr-group-id <integer>] - DR Group ID

If this parameter is specified, the command displays information for the matching DR group.

[-cluster-uuid <UUID>] - Cluster UUID

If this parameter is specified, the command displays information for the matching cluster specified by uuid.

[-cluster <Cluster name>] - Cluster Name

If this parameter is specified, the command displays information for the matching cluster.

[-node-uuid <UUID>] - Node UUID

If this parameter is specified, the command displays information for the matching node specified by uuid.

[-node <text>] - Node Name

If this parameter is specified, the command displays information for the matching nodes.

[-home-port {<netport>|<ifgrp>}] - Home Port

If this parameter is specified, the command displays information for the matching home-port.

[-relationship-type <Roles of MetroCluster Nodes>] - Relationship Role Type

If this parameter is specified, the command displays information for the matching relationship-type.

[-source-address <IP Address>] - Source Network Address

If this parameter is specified, the command displays information for the matching source address.

[-destination-address <IP Address>] - Destination Network Address

If this parameter is specified, the command displays information for the matching destination address.

[-partner-cluster-uuid <UUID>] - Partner Cluster UUID

If this parameter is specified, the command displays information for the matching partner-cluster-uuid.

[-partner-node-uuid <UUID>] - Partner Node UUID

If this parameter is specified, the command displays information for the matching partner-node-uuid.

[-partner-node <text>] - Partner Node Name

If this parameter is specified, the command displays information for the matching partner-node.

[-partner-type <text>] - Partner Relationship Type

If this parameter is specified, the command displays information for the matching partner-type.

[-config-state <text>] - Configuration State

If this parameter is specified, the command displays information for the matching config-state.

[-config-error-info <text>] - Configuration Error Information

If this parameter is specified, the command displays information for the matching config-error-info.

Examples

The following example shows the output of metrocluster configuration-settings connection connect command:

```
Output of the command before the connections are established using the
xref:{relative_path}metrocluster-configuration-settings-connection-
connect.html[metrocluster configuration-settings connection connect]
command:
clusA::> metrocluster configuration-settings connection show
DR                    Source          Destination
Group Cluster Node    Network Address Network Address Partner Type Config
State
----- ------- ------- --------------- --------------- ------------
------------
1     clusA A1
                 Home Port: e0f
                      10.140.113.214  10.140.113.216  HA Partner
disconnected
                 Home Port: e0f
                      10.140.113.214  10.140.113.218  DR Partner
disconnected
                 Home Port: e0f
                      10.140.113.214  10.140.113.249  DR Auxiliary
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.217  HA Partner
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.248  DR Partner
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.25   DR Auxiliary
disconnected
                       A2
                 Home Port: e0f
                      10.140.113.216  10.140.113.214  HA Partner
disconnected
                 Home Port: e0f
```

```
                     10.140.113.216  10.140.113.249  DR Partner
disconnected
                Home Port: e0f
                     10.140.113.216  10.140.113.218  DR Auxiliary
disconnected
                Home Port: e0g
                     10.140.113.217  10.140.113.215  HA Partner
disconnected
                Home Port: e0g
                     10.140.113.217  10.140.113.25   DR Partner
disconnected
                Home Port: e0g
                     10.140.113.217  10.140.113.248  DR Auxiliary
disconnected
      clusB B2
                Home Port: e0f
                     10.140.113.249  10.140.113.218  HA Partner
disconnected
                Home Port: e0f
                     10.140.113.249  10.140.113.216  DR Partner
disconnected
                Home Port: e0f
                     10.140.113.249  10.140.113.214  DR Auxiliary
disconnected
                Home Port: e0g
                     10.140.113.25   10.140.113.248  HA Partner
disconnected
                Home Port: e0g
                     10.140.113.25   10.140.113.217  DR Partner
disconnected
                Home Port: e0g
                     10.140.113.25   10.140.113.215  DR Auxiliary
disconnected
                      B1
                Home Port: e0f
                     10.140.113.218  10.140.113.249  HA Partner
disconnected
                Home Port: e0f
                     10.140.113.218  10.140.113.214  DR Partner
disconnected
                Home Port: e0f
                     10.140.113.218  10.140.113.216  DR Auxiliary
disconnected
                Home Port: e0g
                     10.140.113.248  10.140.113.25   HA Partner
disconnected
```

```
                 Home Port: e0g
                      10.140.113.248  10.140.113.215  DR Partner
disconnected
                 Home Port: e0g
                      10.140.113.248  10.140.113.217  DR Auxiliary
disconnected
24 entries were displayed.
Output of the command after the connections are established using the
xref:{relative_path}metrocluster-configuration-settings-connection-
connect.html[metrocluster configuration-settings connection connect]
command:
clusA::>  metrocluster configuration-settings connection show
DR                    Source          Destination
Group Cluster Node    Network Address Network Address Partner Type Config
State
----- ------- ------- --------------- --------------- ------------
------------
1     clusA A1
                 Home Port: e0f
                      10.140.113.214  10.140.113.216  HA Partner
completed
                 Home Port: e0f
                      10.140.113.214  10.140.113.218  DR Partner
completed
                 Home Port: e0f
                      10.140.113.214  10.140.113.249  DR Auxiliary
completed
                 Home Port: e0g
                      10.140.113.215  10.140.113.217  HA Partner
completed
                 Home Port: e0g
                      10.140.113.215  10.140.113.248  DR Partner
completed
                 Home Port: e0g
                      10.140.113.215  10.140.113.25   DR Auxiliary
completed
                       A2
                 Home Port: e0f
                      10.140.113.216  10.140.113.214  HA Partner
completed
                 Home Port: e0f
                      10.140.113.216  10.140.113.249  DR Partner
completed
                 Home Port: e0f
                      10.140.113.216  10.140.113.218  DR Auxiliary
completed
```

```
                 Home Port: e0g
                      10.140.113.217  10.140.113.215  HA Partner
completed
                 Home Port: e0g
                      10.140.113.217  10.140.113.25   DR Partner
completed
                 Home Port: e0g
                      10.140.113.217  10.140.113.248  DR Auxiliary
completed
      clusB B2
                 Home Port: e0f
                      10.140.113.249  10.140.113.218  HA Partner
completed
                 Home Port: e0f
                      10.140.113.249  10.140.113.216  DR Partner
completed
                 Home Port: e0f
                      10.140.113.249  10.140.113.214  DR Auxiliary
completed
                 Home Port: e0g
                      10.140.113.25   10.140.113.248  HA Partner
completed
                 Home Port: e0g
                      10.140.113.25   10.140.113.217  DR Partner
completed
                 Home Port: e0g
                      10.140.113.25   10.140.113.215  DR Auxiliary
completed
                       B1
                 Home Port: e0f
                      10.140.113.218  10.140.113.249  HA Partner
completed
                 Home Port: e0f
                      10.140.113.218  10.140.113.214  DR Partner
completed
                 Home Port: e0f
                      10.140.113.218  10.140.113.216  DR Auxiliary
completed
                 Home Port: e0g
                      10.140.113.248  10.140.113.25   HA Partner
completed
                 Home Port: e0g
                      10.140.113.248  10.140.113.215  DR Partner
completed
                 Home Port: e0g
                      10.140.113.248  10.140.113.217  DR Auxiliary
```

```
completed
24 entries were displayed.
```

Related Links

- metrocluster configuration-settings connection connect

metrocluster configuration-settings dr-group create

Create a DR group in a MetroCluster over IP setup

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

Description

The `metrocluster configuration-settings dr-group create` command partners the nodes that will make up a DR group in a MetroCluster setup.

This command is used for MetroCluster setups that are connected though IP links. MetroCluster setups that are connected through FC links will configure DR groups automatically and do not require the `metrocluster configuration-settings` commands.

The `metrocluster configuration-settings` commands are run in the following order to set up MetroCluster:

- `metrocluster configuration-settings dr-group create` ,
- metrocluster configuration-settings interface create ,
- metrocluster configuration-settings connection connect .

Before running this command, cluster peering must be configured between the local and partner clusters. Run the cluster peer show command to verify that peering is available between the local and partner clusters.

This command configures a local node and a remote node as DR partner nodes. The command also configures the HA partner of the local node and the HA partner of the remote node as the other DR partner nodes in the DR group.

Parameters

`-partner-cluster <Cluster name>` - Partner Cluster Name

Use this parameter to specify the name of the partner cluster.

`-local-node {<nodename>|local}` - Local Node Name

Use this parameter to specify the name of a node in the local cluster.

`-remote-node <text>` - Remote Node Name

Use this parameter to specify the name of a node in the partner cluster that is to be the DR partner of the specified local node.

Examples

```
The following example shows the creation of the MetroCluster DR group:
```

```
clusA::>  metrocluster configuration-settings dr-group  create -partner
-cluster clusB -local-node A1 -remote-node B1
[Job 268] Job succeeded: DR Group Create is successful.

clusA::>  metrocluster configuration-settings dr-group show
DR Group ID Cluster                    Node                DR Partner Node
----------- -------------------------- ------------------
------------------
1           clusA
                                       A1                  B1
                                       A2                  B2
            clusB
                                       B2                  A2
                                       B1                  A1
4 entries were displayed.
clusA::> metrocluster configuration-settings show-status
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
---------------------------------
clusA
                           A1                  ready for interface create
                           A2                  ready for interface create
clusB
                           B1                  ready for interface create
                           B2                  ready for interface create
4 entries were displayed.
```

Related Links

- metrocluster configuration-settings interface create
- metrocluster configuration-settings connection connect
- cluster peer show

metrocluster configuration-settings dr-group delete

Delete a DR group in a MetroCluster over IP setup

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

Description

The `metrocluster configuration-settings dr-group delete` command deletes a DR group and its node partnerships that were configured using the metrocluster configuration-settings dr-group create command.

This command cannot be run if the metrocluster configuration-settings interface create command has configured a network logical interface on a network port provisioned for MetroCluster. The metrocluster configuration-settings interface delete command must then be run to delete the network logical interfaces on every node in the DR group.

The `metrocluster configuration-settings` commands are run in the following order to remove the MetroCluster over IP configuration:

- metrocluster configuration-settings connection disconnect ,
- metrocluster configuration-settings interface delete ,
- `metrocluster configuration-settings dr-group delete` .

Parameters

`-dr-group-id <integer>` - Dr group Id

This parameter identifies the DR group to be deleted.

Examples

```
The following example shows the deletion of the MetroCluster DR group:
```

```
clusA::> metrocluster configuration-settings dr-group delete -dr-group-id
1

Warning: This command deletes the existing DR group relationship. Are you
sure
         you want to proceed ? {y|n}: y
[Job 279] Job succeeded: DR Group Delete is successful.

clusA::>  metrocluster configuration-settings dr-group  show
No DR groups exist.

clusA::> metrocluster configuration-settings show-status
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
---------------------------------
clusA
                           A1                  ready for DR group create
                           A2                  ready for DR group create
clusB
                           B1                  ready for DR group create
                           B2                  ready for DR group create
4 entries were displayed.
```

Related Links

- metrocluster configuration-settings dr-group create
- metrocluster configuration-settings interface create
- metrocluster configuration-settings interface delete
- metrocluster configuration-settings connection disconnect

metrocluster configuration-settings dr-group show

Display the DR groups in a MetroCluster over IP setup

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

Description

The `metrocluster configuration-settings dr-group show` command displays the DR groups and their nodes in the MetroCluster over IP setup.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you

specify.

| [-instance] }

If this parameter is specified, the command displays detailed information about all entries.

[-dr-group-id <integer>] - DR Group ID

If this parameter is specified, the command displays information for the matching DR group.

[-cluster-uuid <UUID>] - Cluster UUID

If this parameter is specified, the command displays information for the matching cluster uuid.

[-cluster <Cluster name>] - Cluster Name

If this parameter is specified, the command displays information for the specified cluster.

[-node-uuid <UUID>] - Node UUID

If this parameter is specified, the command displays information for the matching nodes uuid.

[-node <text>] - Node Name

If this parameter is specified, the command displays information for the matching nodes.

[-dr-partner-node-uuid <UUID>] - DR Partner Node UUID

If this parameter is specified, the command displays information for the matching DR partner node uuid.

[-dr-partner-node <text>] - DR Partner Node Name

If this parameter is specified, the command displays information for the matching DR partner nodes.

Examples

The following example illustrates the display of DR group configuration in a four-node MetroCluster setup:

```
clusA::>  metrocluster configuration-settings dr-group show
DR Group ID Cluster                    Node                DR Partner Node
----------- -------------------------- ------------------
------------------
1           clusA
                                       A1                  B1
                                       A2                  B2
            clusB
                                       B2                  A2
                                       B1                  A1
4 entries were displayed.
```

metrocluster configuration-settings interface create

Create a MetroCluster interface

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

Description

The `metrocluster configuration-settings interface create` command configures the network logical interfaces that will be used on a node in a MetroCluster setup to mirror NV logs and access remote storage.

This command is used for MetroCluster setups that are connected though IP links. MetroCluster setups that are connected through FC links do not require the user to provision network logical interfaces to mirror NV logs and access remote storage.

The `metrocluster configuration-settings` commands are run in the following order to set up MetroCluster:

- metrocluster configuration-settings dr-group create ,
- `metrocluster configuration-settings interface create` ,
- metrocluster configuration-settings connection connect .

Before running this command , the node's DR group must be configured using the metrocluster configuration-settings dr-group create command. Run the metrocluster configuration-settings dr-group show command to verify that the node's DR group has been configured.

Parameters

`-cluster-name <Cluster name>` - Cluster Name

Use this parameter to specify the name of the local or partner cluster.

`-home-node <text>` - Home Node

Use this parameter to specify the home node in the cluster which hosts the interface.

`-home-port {<netport>|<ifgrp>}` - Home Port

Use this parameter to specify the home port provisioned for MetroCluster.

`-address <IP Address>` - Network Address

Use this parameter to specify the network address to be assigned to the home port.

`-netmask <Contiguous IP Mask>` - Netmask

Use this parameter to specify the network mask to be assigned to the interface.

`[-gateway <IP Address>]` - Gateway

Use this parameter to specify the gateway to be assigned for the routed network.

`[-vlan-id <integer>]` - Virtual LAN ID

Use this parameter to specify the VLAN id.

Examples

This example shows configuring logical interface on MetroCluster IP capable port:

```
clusA::>  metrocluster configuration-settings interface create -cluster
-name clusA -home-node A1 -home-port e0f -address 10.140.113.214 -netmask
```

```
255.255.192.0
[Job 281] Job succeeded: Interface Create is successful.

clusA::> metrocluster configuration-settings interface show
DR
Config
Group Cluster Node    Network Address Netmask         Gateway
State
----- ------- ------- --------------- --------------- ---------------
---------
1     clusA A1
                 Home Port: e0f
                                       10.140.113.214  255.255.192.0   -
completed
Output after configuring all the interfaces:
clusA::> metrocluster configuration-settings interface show
DR
Config
Group Cluster Node    Network Address Netmask         Gateway
State
----- ------- ------- --------------- --------------- ---------------
---------
1     clusA A1
                 Home Port: e0f
                      10.140.113.214  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.215  255.255.192.0   -
completed
                       A2
                 Home Port: e0f
                      10.140.113.216  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.217  255.255.192.0   -
completed
      clusB B2
                 Home Port: e0f
                      10.140.113.249  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.25   255.255.192.0   -
completed
                       B1
                 Home Port: e0f
                      10.140.113.218  255.255.192.0   -
```

```
completed
                 Home Port: e0g
                      10.140.113.248  255.255.192.0   -
completed
8 entries were displayed.

clusA::> metrocluster configuration-settings show-status
Cluster                    Node               Configuration Settings
Status
-------------------------- ------------------
---------------------------------
clusA
                           A1                 ready for connection connect
                           A2                 ready for connection connect
clusB
                           B1                 ready for connection connect
                           B2                 ready for connection connect
4 entries were displayed.

clusA::> metrocluster configuration-settings connection show
DR                     Source          Destination
Group Cluster Node    Network Address Network Address Partner Type Config
State
----- ------- ------- --------------- --------------- ------------
------------
1     clusA A1
                 Home Port: e0f
                      10.140.113.214  10.140.113.216  HA Partner
disconnected
                 Home Port: e0f
                      10.140.113.214  10.140.113.218  DR Partner
disconnected
                 Home Port: e0f
                      10.140.113.214  10.140.113.249  DR Auxiliary
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.217  HA Partner
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.248  DR Partner
disconnected
                 Home Port: e0g
                      10.140.113.215  10.140.113.25   DR Auxiliary
disconnected
                       A2
                 Home Port: e0f
```

```
                     10.140.113.216  10.140.113.214  HA Partner
disconnected
                Home Port: e0f
                     10.140.113.216  10.140.113.249  DR Partner
disconnected
                Home Port: e0f
                     10.140.113.216  10.140.113.218  DR Auxiliary
disconnected
                Home Port: e0g
                     10.140.113.217  10.140.113.215  HA Partner
disconnected
                Home Port: e0g
                     10.140.113.217  10.140.113.25   DR Partner
disconnected
                Home Port: e0g
                     10.140.113.217  10.140.113.248  DR Auxiliary
disconnected
      clusB B2
                Home Port: e0f
                     10.140.113.249  10.140.113.218  HA Partner
disconnected
                Home Port: e0f
                     10.140.113.249  10.140.113.216  DR Partner
disconnected
                Home Port: e0f
                     10.140.113.249  10.140.113.214  DR Auxiliary
disconnected
                Home Port: e0g
                     10.140.113.25   10.140.113.248  HA Partner
disconnected
                Home Port: e0g
                     10.140.113.25   10.140.113.217  DR Partner
disconnected
                Home Port: e0g
                     10.140.113.25   10.140.113.215  DR Auxiliary
disconnected
                      B1
                Home Port: e0f
                     10.140.113.218  10.140.113.249  HA Partner
disconnected
                Home Port: e0f
                     10.140.113.218  10.140.113.214  DR Partner
disconnected
                Home Port: e0f
                     10.140.113.218  10.140.113.216  DR Auxiliary
disconnected
```

```
                  Home Port: e0g
                       10.140.113.248  10.140.113.25   HA Partner
disconnected
                  Home Port: e0g
                       10.140.113.248  10.140.113.215  DR Partner
disconnected
                  Home Port: e0g
                       10.140.113.248  10.140.113.217  DR Auxiliary
disconnected
24 entries were displayed.
```

Related Links

- metrocluster configuration-settings dr-group create
- metrocluster configuration-settings connection connect
- metrocluster configuration-settings dr-group show

metrocluster configuration-settings interface delete

Delete a MetroCluster interface

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

Description

The `metrocluster configuration-settings interface delete` command deletes the network logical interface that was configured on a network port provisioned for MetroCluster.

This command cannot be run if the metrocluster configuration-settings connection connect command has set up the connections between the nodes in a DR group. The metrocluster configuration-settings connection disconnect command must then be run to remove the connections.

The `metrocluster configuration-settings` commands are run in the following order to remove the MetroCluster over IP configuration:

- metrocluster configuration-settings connection disconnect ,
- `metrocluster configuration-settings interface delete` ,
- metrocluster configuration-settings dr-group delete .

Parameters

`-cluster-name <Cluster name>` - Cluster Name

Use this parameter to specify the name of the local or partner cluster.

`-home-node <text>` - Home Node

Use this parameter to specify the home node in the cluster which hosts the interface.

-home-port {<netport>|<ifgrp>} - Home Port

Use this parameter to specify the home port provisioned for MetroCluster.

Examples

The following example shows the deletion of interface in a MetroCluster setup:

```
clusA::> metrocluster configuration-settings interface delete -cluster
-name clusA -home-node A1 -home-port e0f
[Job 271] Job succeeded: Interface Delete is successful.

clusA::> metrocluster configuration-settings interface show
DR
Config
Group Cluster Node    Network Address Netmask         Gateway
State
----- ------- ------- --------------- --------------- ---------------
---------
1     clusA A1
                 Home Port: e0g
                      10.140.113.215  255.255.192.0   -
completed
                       A2
                 Home Port: e0f
                      10.140.113.216  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.217  255.255.192.0   -
completed
      clusB B2
                 Home Port: e0f
                      10.140.113.249  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.25   255.255.192.0   -
completed
                       B1
                 Home Port: e0f
                      10.140.113.218  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.248  255.255.192.0   -
completed
7 entries were displayed.

clusA::> metrocluster configuration-settings show-status
```

```
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
---------------------------------
clusA
                           A1                  ready for next interface
create
                           A2                  ready for connection connect
clusB
                           B1                  ready for connection connect
                           B2                  ready for connection connect
4 entries were displayed.
Output of the command after deleting all the interfaces:
clusA::> metrocluster configuration-settings interface show
No interfaces exist.

clusA::> metrocluster configuration-settings show-status
Cluster                    Node                Configuration Settings
Status
-------------------------- ------------------
---------------------------------
clusA
                           A1                  ready for interface create
                           A2                  ready for interface create
clusB
                           B1                  ready for interface create
                           B2                  ready for interface create
4 entries were displayed.
```

Related Links

- metrocluster configuration-settings connection connect
- metrocluster configuration-settings connection disconnect
- metrocluster configuration-settings dr-group delete

metrocluster configuration-settings interface show

Display the network logical interfaces provisioned for MetroCluster

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

Description

The `metrocluster configuration-settings interface show` command displays the network logical interfaces that were provisioned for MetroCluster.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| [-instance] }

If this parameter is specified, the command displays detailed information about all entries.

[-dr-group-id <integer>] - DR Group ID

If this parameter is specified, the command displays information for the matching DR group.

[-cluster-uuid <UUID>] - Cluster UUID

If this parameter is specified, the command displays information for the matching cluster specified by uuid.

[-cluster <Cluster name>] - Cluster Name

If this parameter is specified, the command displays information for the matching cluster..

[-node-uuid <UUID>] - Node UUID

If this parameter is specified, the command displays information for the matching nodes uuid.

[-node <text>] - Node Name

If this parameter is specified, the command displays information for the matching nodes.

[-home-port {<netport>|<ifgrp>}] - Home Port

If this parameter is specified, all interfaces with home-port set to this value are displayed.

[-address <IP Address>] - Network Address

If this parameter is specified, the command displays information for the matching network address.

[-netmask <Contiguous IP Mask>] - Netmask

If this parameter is specified, all interfaces with netmask set to this value are displayed.

[-gateway <IP Address>] - Gateway

If this parameter is specified, all interfaces with gateway set to this value are displayed.

[-config-state <text>] - Configuration State

If this parameter is specified, all interfaces with this field set to the specified value are displayed.

[-config-error-info <text>] - Configuration Error Information

If this parameter is specified, all interfaces with this field set to the specified value are displayed.

[-vlan-id <integer>] - Virtual LAN ID

If this parameter is specified, all interfaces with vlan-id set to this value are displayed.

Examples

The following example illustrates display of logical interfaces configured in a four-node MetroCluster setup:

```
clusA::> metrocluster configuration-settings interface show
DR
Config
Group Cluster Node    Network Address Netmask         Gateway
State
----- ------- ------- --------------- --------------- ---------------
---------
1     clusA A1
                 Home Port: e0f
                      10.140.113.214  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.215  255.255.192.0   -
completed
                       A2
                 Home Port: e0f
                      10.140.113.216  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.217  255.255.192.0   -
completed
      clusB B2
                 Home Port: e0f
                      10.140.113.249  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.25   255.255.192.0   -
completed
                       B1
                 Home Port: e0f
                      10.140.113.218  255.255.192.0   -
completed
                 Home Port: e0g
                      10.140.113.248  255.255.192.0   -
completed
8 entries were displayed.
```

metrocluster configuration-settings mediator add

Configure the network connections between the Mediator and MetroCluster nodes

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

Description

The `metrocluster configuration-settings mediator add` command configures the connection between the mediator and all the nodes in a MetroCluster configuration.

Before this command is run MetroCluster should be configured on all the clusters.

Mediator username and password are required to successfully add the mediator.

After this command completes successfully, every node will:

- Have a connection with the mediator.
- The mediator disks will be assigned to the nodes in the MetroCluster configuration.
- AUSO will be enabled.

Parameters

`-mediator-address <IP Address>` - Mediator IP Address

Specifies the IP address of the mediator to which the nodes in the MetroCluster configuration will be connected.

Examples

```
The following example connects the nodes to the mediator in a MetroCluster
configuration over IP setup:
```

```
clusA::> metrocluster configuration-settings mediator add -mediator
-address 10.234.133.115
        Adding mediator and enabling Automatic Unplanned Switchover. It
might take a few minutes to complete.
        Please enter the username for the mediator: mediatoradmin
        Please enter the password for the mediator:
        Confirm the mediator password:
        Creating mediator mailboxes...
        Setting up connections to mediator from all nodes in the
clusters...
        Setting mediator mailbox from all nodes in the cluster...
        Enabling Automatic Unplanned Switchover for all nodes in the
cluster...
        Successfully added mediator.
```

metrocluster configuration-settings mediator remove

Tear down connections between the Mediator and MetroCluster nodes

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

Description

The `metrocluster configuration-settings mediator remove` command removes the connection between the mediator and all the nodes in a MetroCluster configuration.

Mediator username and password are required to successfully remove the mediator.

Parameters

`[-dr-group-id <integer>]` - DR Group Id (privilege: advanced)

Specifies the Disaster Recovery Group Identifier for which the mediator connections need to be removed.

Examples

```
The following example removes the connections between the nodes and the
mediator in a MetroCluster configuration over IP setup:
```

```
clusA::> metrocluster configuration-settings mediator remove
      Removing the mediator and disabling Automatic Unplanned Switchover.
It may take a few minutes to complete.
      Please enter the username for the mediator: mediatoradmin
      Please enter the password for the mediator:
      Confirm the mediator password:
      Disabling Automatic Unplanned Switchover for all nodes in the
cluster...
      Removing mediator mailboxes...
      Performing final cleanup...
      Successfully removed the mediator.
```

metrocluster configuration-settings mediator show

Display the nodes connected to the mediator

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

Description

The `metrocluster configuration-settings mediator show` command displays the connection status of the nodes with the Mediator in a MetroCluster configuration.

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.

[-mediator-address <IP Address>] - Mediator IP Address

Specifies the IP address of the mediator to which the nodes in the MetroCluster configuration are connected.

[-port <integer>] - Mediator port number

Specifies the port of the mediator to which the nodes in the MetroCluster configuration are connected.

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

Specifies the nodes in the MetroCluster configuration which are connected to the mediator.

[-reachable {true|false}] - Connection Status of the Mediator

Specifies the connection status of the nodes with the mediator in the MetroCluster configuration.

[-configured {true|false}] - Mediator Configuration Status

Specifies the configuration status of the nodes with the mediator in the MetroCluster configuration.

Examples

```
The following example shows the mediator connection status in a
MetroCluster configuration over IP setup:
```

```
cluster1_node_01::*> metrocluster configuration-settings mediator show
        Mediator IP     Port    Node                    Configuration
Connection
                                                        Status
Status
        --------------- ------- ----------------------- -------------
-----------
        10.234.217.168
                        31784   cluster1_node_01        true          true
                                cluster1_node_02        true          true
                                cluster2_node_01        true          true
                                cluster2_node_02        true          true
```

metrocluster interconnect commands

metrocluster interconnect adapter modify

Modify MetroCluster interconnect adapter settings

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

Description

The `metrocluster interconnect adapter modify` command enables you to modify settings of the MetroCluster interconnect adapter.

Parameters

`-node {<nodename>|local}` - Node Name (privilege: advanced)

This parameter specifies the node name.

`-is-ood-enabled {true|false}` - Is Out-of-Order Delivery Enabled? (privilege: advanced)

This parameter specifies the out-of-order delivery setting on the adapter.

Examples

The following example enables out-of-order delivery for the port 'fcvi_device_0' on the node 'clusA-01':

```
clusA::*> metrocluster interconnect adapter modify -node clusA-01 -adapter
-port-name fcvi_device_0 -is-ood-enabled true
```

metrocluster interconnect adapter show

Display MetroCluster interconnect adapter information

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

Description

The `metrocluster interconnect adapter show` command displays interconnect adapter information for the nodes in a MetroCluster configuration.

This command displays the following details about the local node and the HA partner node:

- Node: This field specifies the name of the node in the cluster.
- Adapter Name: This field specifies the name of the interconnect adapter.
- Adapter Type: This field specifies the type of the interconnect adapter.
- Link Status: This field specifies the physical link status of the interconnect adapter.
- Is OOD Enabled: This field specifies the out-of-order delivery status of the interconnect adapter.
- IP Address: This field specifies the IP address assigned to the interconnect adapter.
- Port Number: This field specifies the port number of the interconnect adapter.

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.

| [-connectivity]

Displays the connectivity information from all the interconnect adapters to the connected nodes.

| [-connectivity-hidden] (privilege: advanced)

Displays additional connectivity information (IP address, Area ID, Port ID) from all the interconnect adapters to the connected nodes.

| [-switch]

Displays details of switches connected to all the interconnect adapters.

| [-instance] }

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

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

Displays information only about the interconnect adapters that are hosted by the specified node.

[-adapter <text>] - Adapter

Displays information only about the interconnect adapters that match the specified name.

[-port-name <text>] - Port Name

Displays information only about the interconnect adapters that host the specified port name.

[-type <text>] - Adapter Type

Displays information only about the interconnect adapters that match the specified adapter type.

[-physical-status <text>] - Physical Status

Displays information only about the interconnect adapters that match the specified physical status.

[-wwn <text>] - Adapter Port World Wide Name

Displays information only about the interconnect adapters that match the specified world wide name.

[-address <text>] - IP Address

Displays information only about the interconnect adapters that match the specified IP address.

[-firmware-version <text>] - Firmware Version

Displays information only about the interconnect adapters that match the specified firmware version.

[-link-speed <text>] - Link Speed

Displays information only about the interconnect adapters that match the specified link speed.

[-link-speed-neg-type <text>] - Link Speed Negotiation Type

Displays information only about the interconnect adapters that match the specified negotiated link speed type.

[-switch-name <text>] - Switch Name

Displays information only about the interconnect adapters that are connected to the specified switch.

[-switch-model <text>] - Switch Model

Displays information only about the interconnect adapters that are connected to the switch with the specified model.

[-switch-wwn <text>] - Switch WWName

Displays information only about the interconnect adapters that are connected to the switch with the specified world wide name.

[-switch-vendor <text>] - Switch Vendor

Displays information only about the interconnect adapters that are connected to the switch with the specified vendor.

[-switch-status <text>] - Switch Status

Displays information only about the interconnect adapters that are connected to the switch with the specified operational status.

[-switch-port-number <text>] - Switch Port Number

Displays information only about the interconnect adapters that are connected to the switch with the specified port number.

[-switch-port-wwpn <text>] - Switch Port WWPN

Displays information only about the interconnect adapters that are connected to the switch with the specified word wide port name.

[-remote-adapter-name-list <text>,…] - Remote Adapter Name List

Displays information only about the interconnect adapters that are connected to the specified remote adapters.

[-remote-adapter-wwn-list <text>,…] - Remote Adapter WWName List

Displays information only about the interconnect adapters that are connected to the remote adapters with the specified world wide names.

[-remote-adapter-address-list <text>,…] - Remote Adapter IP Address List

Displays information only about the interconnect adapters that are connected to the remote adapters with the specified IP addresses.

[-remote-adapter-port-id-list <Hex Integer>,…] - Remote Adapter Port ID List

Displays information only about the interconnect adapters that are connected to the remote adapters with the specified port IDs.

[-remote-adapter-domain-id-list <integer>,…] - Remote Adapter Domain ID List

Displays information only about the interconnect adapters that are connected to the remote adapters with the specified domain IDs.

[-remote-adapter-area-id-list <integer>,…] - Remote Adapter Area ID List

Displays information only about the interconnect adapters that are connected to the remote adapters with the specified Area IDs.

[-remote-partner-system-id-list <integer>,…] - Remote Partner System ID List

Displays information only about the interconnect adapters that are connected to the remote nodes with the

specified System IDs.

[-remote-partner-name-list {<nodename>|local}] - Remote Partner Name List

Displays information only about the interconnect adapters that are connected to the specified remote nodes.

[-is-ood-enabled {true|false}] - Is Out-of-Order Delivery Enabled?

Displays information only about the interconnect adapters that match the specified out-of-order delivery setting.

Examples

The following example shows the output of the command during normal operation (neither cluster is in switchover state):

```
clusA::> metrocluster interconnect adapter show
                                          Link   Is OOD
       Node         Adapter         Type  Status Enabled? IP Address  Port
Number
       ------------ --------------- ----- ------ -------- -----------
-----------
       clusA-01     cxgb3_0         iWARP Up     false    10.0.1.1    c0a
       clusA-01     cxgb3_0         iWARP Down   false    10.0.2.1    c0b
       clusA-01     fcvi_device_0   FC-VI Up     false    1.0.0.1     1a
       clusA-01     fcvi_device_1   FC-VI Up     false    2.0.0.3     1b
       clusA-02     cxgb3_0         iWARP Up     false    10.0.1.2    c0a
       clusA-02     cxgb3_0         iWARP Down   false    10.0.2.2    c0b
       clusA-02     fcvi_device_0   FC-VI Up     false    1.0.1.1     1a
       clusA-02     fcvi_device_1   FC-VI Up     false    2.0.1.3     1b
```

The following example shows the output of the command after MetroCluster switchover is performed:

```
clusA::> metrocluster interconnect adapter show
                                          Link   Is OOD
       Node         Adapter         Type  Status Enabled? IP Address  Port
Number
       ------------ --------------- ----- ------ -------- -----------
-----------
       clusA-01     cxgb3_0         iWARP Up     false    10.0.1.1    c0a
       clusA-01     cxgb3_0         iWARP Down   false    10.0.2.1    c0b
       clusA-01     fcvi_device_0   FC-VI Down   false    1.0.0.1     1a
       clusA-01     fcvi_device_1   FC-VI Down   false    2.0.0.3     1b
       clusA-02     cxgb3_0         iWARP Up     false    10.0.1.2    c0a
       clusA-02     cxgb3_0         iWARP Down   false    10.0.2.2    c0b
       clusA-02     fcvi_device_0   FC-VI Down   false    1.0.1.1     1a
       clusA-02     fcvi_device_1   FC-VI Down   false    2.0.1.3     1b
```

The following example shows the output of the command with connectivity field during normal operation (neither cluster is in switchover state):

```
clusA::> metrocluster interconnect adapter show -connectivity -node local
-type FC-VI
Adapter Name: fcvi_device_0
                     WWName: 21:00:00:24:ff:32:01:68
                     PortNo: 1a
Remote Adapters:
Adapter Name Partner Node Name World Wide Name         PortId
       ------------ ----------------- ----------------------- ------
       fcvi_device_0
                    clusA-01          21:00:00:24:ff:32:01:80  65536
       fcvi_device_0
                    clusB-01          21:00:00:24:ff:32:01:54 131072
       fcvi_device_0
                    clusB-02          21:00:00:24:ff:32:01:60 131328
Adapter Name: fcvi_device_1
                     WWName: 21:00:00:24:ff:32:01:69
                     PortNo: 1b
Remote Adapters:
Adapter Name Partner Node Name World Wide Name         PortId
       ------------ ----------------- ----------------------- ------
       fcvi_device_1
                    clusA-01          21:00:00:24:ff:32:01:81 196608
       fcvi_device_1
                    clusB-01          21:00:00:24:ff:32:01:55 262144
       fcvi_device_1
                    clusB-02          21:00:00:24:ff:32:01:61 262400
```

The following example shows the output of the command with connectivity field after MetroCluster switchover is performed.

```
clusA::> metrocluster interconnect adapter show -connectivity -node local
-type FC-VI
Adapter Name: fcvi_device_0
                     WWName: 21:00:00:24:ff:32:01:68
                     PortNo: 1a
Remote Adapters:
Adapter Name Partner Node Name World Wide Name         PortId
       ------------ ----------------- ----------------------- ------
       fcvi_device_0
                    clusA-01          21:00:00:24:ff:32:01:80  65536
Adapter Name: fcvi_device_1
                     WWName: 21:00:00:24:ff:32:01:69
                     PortNo: 1b
Remote Adapters:
Adapter Name Partner Node Name World Wide Name         PortId
       ------------ ----------------- ----------------------- ------
       fcvi_device_1
                    clusA-01          21:00:00:24:ff:32:01:81 196608
```

metrocluster interconnect mirror show

Display MetroCluster interconnect mirror information

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

Description

The `metrocluster interconnect mirror show` command displays NVRAM mirror information for the nodes configured in a MetroCluster.

This command displays the following details about the local node and the HA partner node:

- Node: This field specifies the name of the node in the cluster.
- Partner Name: This field specifies the name of the partner node.
- Partner Type: This field specifies the type of the partner.
- Mirror Admin Status: This field specifies the administrative status of the NVRAM mirror between partner nodes.
- Mirror Oper Status: This field specifies the operational status of the NVRAM mirror between partner nodes.
- Adapter: This field specifies the name of the interconnect adapter used for NVRAM mirroring.
- Type: This field specifies the type of the interconnect adapter used for NVRAM mirroring.
- Status: This field specifies the physical status of the interconnect adapter used for NVRAM mirroring.

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>|local}] - Node Name

If this parameter is specified, mirror details of the specified node are displayed.

[-partner-type {HA|DR|AUX}] - Partner Type

If this parameter is specified, mirror details of the specified partner type are displayed.

[-adapter <text>] - Adapter

If this parameter is specified, mirror details of the specified adapter are displayed.

[-type <text>] - Adapter Type

If this parameter is specified, mirror details of the specified adapter type are displayed.

[-status <text>] - Status

If this parameter is specified, mirror details of the adapter with the specified status are displayed.

[-mirror-oper-status {unknown|online|offline}] - Mirror Operational Status

If this parameter is specified, only mirror details with the specified operational status are displayed.

[-partner-name <text>] - Partner Name

If this parameter is specified, mirror details of the specified partner are displayed.

[-mirror-admin-status {enabled|disabled}] - Mirror Administrative Status

If this parameter is specified, only mirror details with the specified administrative status are displayed.

Examples

The following example shows the output of the command during normal operation (neither cluster is in switchover state):

```
clusA::> metrocluster interconnect mirror show
                                Mirror   Mirror
                       Partner Admin    Oper
       Node Partner Name Type    Status   Status  Adapter Type   Status
       ---- ------------ ------- -------- ------- ------- ------ ------
       clusA-01
            clusA-02
                       HA      enabled  online
                                                 cxgb3_0 iWARP  Up
                                                 cxgb3_0 iWARP  Up
            clusB-01
                       DR      enabled  online
                                                 fcvi_device_0
                                                         FC-VI  Up
                                                 fcvi_device_1
                                                         FC-VI  Up
       clusA-02
            clusA-01
                       HA      enabled  online
                                                 cxgb3_0 iWARP  Up
                                                 cxgb3_0 iWARP  Up
            clusB-02
                       DR      enabled  online
                                                 fcvi_device_0
                                                         FC-VI  Up
                                                 fcvi_device_1
                                                         FC-VI  Up
```

The following example shows the output of the command after MetroCluster switchover is performed:

```
clusA::> metrocluster interconnect mirror show
                                Mirror   Mirror
                       Partner Admin    Oper
       Node Partner Name Type    Status   Status  Adapter Type   Status
       ---- ------------ ------- -------- ------- ------- ------ ------
       clusA-01
            clusA-02
                       HA      enabled  online
                                                  cxgb3_0 iWARP  Up
                                                  cxgb3_0 iWARP  Up
            clusB-01
                       DR      disabled offline
                                                  fcvi_device_0
                                                          FC-VI  Up
                                                  fcvi_device_1
                                                          FC-VI  Up
       clusA-02
            clusA-01
                       HA      enabled  online
                                                  cxgb3_0 iWARP  Up
                                                  cxgb3_0 iWARP  Up
            clusB-02
                       DR      disabled offline
                                                  fcvi_device_0
                                                          FC-VI  Up
                                                  fcvi_device_1
                                                          FC-VI  Up
```

metrocluster interconnect mirror multipath show

Display multipath information

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

Description

The `metrocluster interconnect mirror multipath show` command displays the NVRAM mirror multipath policy for the nodes configured in a MetroCluster.

This command displays the following details about the local node and the HA partner node:

- Node: This field specifies the name of the node in the cluster.
- Multipath Policy: This field specifies the multipath policy used for NVRAM mirroring.

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>|local}] - Node Name

If this parameter is specified, mirror details of the specified node are displayed.

[-multipath-policy {no-mp|static-map|dynamic-map|round-robin}] - Multipath Policy

If this parameter is specified, nodes with the specified multipath policy are displayed.

Examples

The following example shows the output of the command:

```
clusA::> metrocluster interconnect mirror multipath show
       Node            Multipath Policy
       --------------- -----------------
       clusA-1         static-map
       clusA-2         static-map
```

metrocluster node commands

metrocluster node show

Display MetroCluster node configuration information

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

Description

The `metrocluster node show` command displays configuration information for the nodes in the MetroCluster configuration.

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.

| [-partners]

If this option is used the MetroCluster node partnership view will be displayed.

| [-instance] }

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

[-dr-group-id <integer>] - DR Group ID

If this parameter is specified, all nodes belonging to the specified DR group are displayed.

[-cluster <Cluster name>] - Cluster Name

If this parameter is specified, all nodes belonging to the specified cluster are displayed.

[-node <Node name>] - Node Name

If this parameter is specified, the specified node is displayed.

[-ha-partner <Node name>] - HA Partner Name

If this parameter is specified, the node with the specified HA partner is displayed.

[-dr-cluster <Cluster name>] - DR Cluster Name

If this parameter is specified, all nodes belonging to the specified cluster are displayed.

[-dr-partner <Node name>] - DR Partner Name

If this parameter is specified, the node with the specified DR partner is displayed.

[-dr-auxiliary <Node name>] - DR Auxiliary Name

If this parameter is specified, the node with the specified DR auxiliary partner is displayed.

[-node-uuid <UUID>] - Node UUID

If this parameter is specified, the node with the specified Uuid is displayed.

[-ha-partner-uuid <UUID>] - HA Partner UUID

If this parameter is specified, the nodes with the specified HA partner is displayed.

[-dr-partner-uuid <UUID>] - DR Partner UUID

If this parameter is specified, the node with the specified DR partner is displayed.

[-dr-auxiliary-uuid <UUID>] - DR Auxiliary UUID

If this parameter is specified, the node with the specified DR auxiliary partner is displayed.

[-node-cluster-uuid <UUID>] - Node Cluster UUID

If this parameter is specified, all nodes belonging to the specified cluster are displayed.

[-ha-partner-cluster-uuid <UUID>] - HA Partner Cluster UUID

If this parameter is specified, all nodes whose HA partner belong to the specified cluster are displayed.

[-dr-partner-cluster-uuid <UUID>] - DR Partner Cluster UUID

If this parameter is specified, all nodes whose DR partner belong to the specified cluster are displayed.

[-dr-auxiliary-cluster-uuid <UUID>] - DR Auxiliary Cluster UUID

If this parameter is specified, all nodes whose DR auxiliary partner belong to the specified cluster are displayed.

[-node-systemid <integer>] - Node System ID

If this parameter is specified, all nodes with the specified system ID are displayed.

[-ha-partner-systemid <integer>] - HA Partner System ID

If this parameter is specified, all nodes with an HA partner with the specified system ID are displayed.

[-dr-partner-systemid <integer>] - DR Partner System ID

If this parameter is specified, all nodes with a DR partner with the specified system ID are displayed.

[-dr-auxiliary-systemid <integer>] - DR Auxiliary System ID

If this parameter is specified, all nodes with a DR auxiliary partner with the specified system ID are displayed.

[-dr-mirroring-state <text>] - State of DR Mirroring Config

If this parameter is specified, all nodes with this field set to the specified value are displayed. This field specifies if the NVRAM mirroring to the DR partner is enabled through the metrocluster configure command. This field needs to be set to "enabled" for the DR mirroring to be active.

[-configuration-state <text>] - Configuration State of Node

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-additional-configuration-info <text>] - Additional Configuration Info

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-dr-operation-state <text>] - DR Operation State

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-dr-operation-time <integer>] - Time to Complete Operation (secs)

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-node-object-limit {on|off}] - Specifies if the Node Object Limits are Enforced

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-node-ha-partner <text>] - Node and its HA Partner

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-automatic-uso {true|false}] - Automatic USO (privilege: advanced)

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-mccip-node {true|false}] - MCC-IP Node (privilege: advanced)

If this parameter is specified, all nodes with this field set to the specified value are displayed.

[-is-encryption-enabled {true|false}] - Encryption Enabled

If this parameter is specified, all nodes with this field set to the specified value are displayed.

Examples

The following example shows the output of the command before the MetroCluster configuration is done:

```
clusA::> metrocluster node show
       DR                              Configuration  DR
       Group Cluster Node              State          Mirroring Mode
       ----- ------- ------------------ -------------- ---------
--------------------
       -     clusA   clusA-01           ready to configure
                                                     -         -
                     clusA-02           ready to configure
                                                     -         -
                     clusA-03           ready to configure
                                                     -         -
                     clusA-04           ready to configure
                                                     -         -
       4 entries were displayed.
clusA::> metrocluster node show -partners
                             Node (HA Partner) DR Partner (DR Auxiliary)
       ----------------------------------------
----------------------------------------
       Cluster:                            clusA -
                                 clusA-01 (-) - (-)
                                 clusA-02 (-) - (-)
                                 clusA-03 (-) - (-)
                                 clusA-04 (-) - (-)
       4 entries were displayed.
```

The following example shows the output of the command when some DR groups in the MetroCluster configuration are not yet configured:

```
clusA::> metrocluster node show
       DR                              Configuration  DR
       Group Cluster Node              State          Mirroring Mode
       ----- ------- ------------------ -------------- ---------
--------------------
       -     clusA   clusA-03           ready to configure
                                                        -         -
                     clusA-04           ready to configure
                                           -         -
       1     clusA   clusA-01           configured     enabled   normal
                     clusA-02           configured     enabled   normal
             clusB   clusB-01           configured     enabled   normal
                     clusB-02           configured     enabled   normal
       6 entries were displayed.
clusA::> metrocluster node show -partners
                              Node (HA Partner) DR Partner (DR Auxiliary)
       ----------------------------------------
----------------------------------------
       Cluster:                              clusA -
                                   clusA-03 (-) - (-)
                                   clusA-04 (-) - (-)
       Cluster:                              clusA clusB
                             clusA-01 (clusA-02) clusB-01 (clusB-02)
                             clusA-02 (clusA-01) clusB-02 (clusB-01)
       Cluster:                              clusB clusA
                             clusB-01 (clusB-02) clusA-01 (clusA-02)
                             clusB-02 (clusB-01) clusA-02 (clusA-01)
       6 entries were displayed.
```

The following example shows the output of the command after after all DR groups in the MetroCluster configuration are configured:

```
clusA::> metrocluster node show
       DR                              Configuration  DR
       Group Cluster Node              State          Mirroring Mode
       ----- ------- ------------------ -------------- ---------
--------------------
       1     clusA   clusA-01           configured     enabled   normal
                     clusA-02           configured     enabled   normal
             clusB   clusB-01           configured     enabled   normal
                     clusB-02           configured     enabled   normal
       2     clusA   clusA-03           configured     enabled   normal
                     clusA-04           configured     enabled   normal
             clusB   clusB-03           configured     enabled   normal
                     clusB-04           configured     enabled   normal
       8 entries were displayed.
clusA::> metrocluster node show -partners
                                Node (HA Partner) DR Partner (DR Auxiliary)
       ----------------------------------------
----------------------------------------
       Cluster:                                clusA clusB
                             clusA-01 (clusA-02) clusB-01 (clusB-02)
                             clusA-02 (clusA-01) clusB-02 (clusB-01)
       Cluster:                                clusB clusA
                             clusB-01 (clusB-02) clusA-01 (clusA-02)
                             clusB-02 (clusB-01) clusA-02 (clusA-01)
       Cluster:                                clusA clusB
                             clusA-03 (clusA-04) clusB-03 (clusB-04)
                             clusA-04 (clusA-03) clusB-04 (clusB-03)
       Cluster:                                clusB clusA
                             clusB-03 (clusB-04) clusA-03 (clusA-04)
                             clusB-04 (clusB-03) clusA-04 (clusA-03)
       8 entries were displayed.
```

Related Links

- metrocluster configure

metrocluster operation commands

metrocluster operation show

Display details of the last MetroCluster operation

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

Description

The `metrocluster operation show` command displays information about the most recent MetroCluster operation run on the local cluster.

This command will display information about all MetroCluster commands except for the commands in the `metrocluster check` directory. This command will not display any information after MetroCluster has been completely unconfigured using the `metrocluster unconfigure` command.

Examples

The following example shows the output of `metrocluster operation show` after running a metrocluster configure command was successful:

```
clusA::> metrocluster operation show
              Operation: configure
                  State: successful
             Start time: 2/15/2013 18:22:46
               End time: 2/15/2013 18:25:18
                 Errors: -
```

Related Links

- metrocluster configure

metrocluster operation history show

Display details of all MetroCluster operations

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

Description

The `metrocluster operation history show` command displays information about all the MetroCluster operations run on the local cluster.

This command will display information about all MetroCluster commands except for the commands in the `metrocluster check` directory. This command will not display any information after MetroCluster has been completely unconfigured using the `metrocluster unconfigure` command.

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.

[-operation-uuid <UUID>] - Identifier for the Operation

This is the UUID of the operation. If this parameter is specified, only the operation with this UUID is displayed.

[-cluster <Cluster name>] - Cluster Where the Command Was Run

This is the name of the cluster where the command was run. If this parameter is specified, only the operations that were run in this cluster are displayed.

[-node-name <Node name>] - Node Where the Command Was run

This is the name of the node where the command was run. If this parameter is specified, only the operations that were run on this node are displayed.

[-operation <MetroCluster Operation Name>] - Name of the Operation

This is the name of the operation. If this parameter is specified, only the operations with this name are displayed.

[-start-time <MM/DD/YYYY HH:MM:SS>] - Start Time

This is the time the operation started execution. If this parameter is specified, only the operations that were started at this time are displayed.

[-state <MetroCluster Operation state>] - State of the Operation

This is the state of the operation. If this parameter is specified, only the operations that are in this state are displayed.

[-end-time <MM/DD/YYYY HH:MM:SS>] - End Time

This is the time the operation completed. If this parameter is specified, only the operations that completed at this time are displayed.

[-error-list <text>,…] - Error List For the Operation

This is the list of errors that were encountered during an operation's execution. If this parameter is specified, only the operations that have the matching errors are displayed.

[-job-id <integer>] - Identifier for the Job

This is the job id for the operation. If this parameter is specified, only the operation that has the matching job id displayed.

[-additional-info <text>] - Additional Info for Auto Heal

This is the completion status of the auto heal aggregates and auto heal root aggregates phases when processing switchover with auto heal.

Examples

The following example shows the output of `metrocluster operation history show` after some MetroCluster operations have been performed:

```
clusA::> metrocluster operation history show
  Operation        State          Start time       End time
  ---------------- -------------- ---------------- ------------------
  configure        successful     2/15/2013 18:22:46
                     2/15/2013 18:25:18
  configure        failed         2/15/2013 18:13:45
                     2/15/2013 18:13:45
  2 entries were displayed.
```

metrocluster transition commands

metrocluster transition disable

Disable Transition

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

Description

The `metrocluster transition disable` command is used to end the transition from a MetroCluster FC to a MetroCluster IP configuration. The command clears the setting for transition mode on the local and partner clusters.

Examples

The following example shows how to disable the transition mode:

```
clusA::> metrocluster transition disable
clusA::> metrocluster transition show-mode
 Transition Mode
 --------------------------
 not-enabled
```

metrocluster transition enable

Enable Transition

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

Description

The `metrocluster transition enable` command is used to start the transition from a MetroCluster FC to a MetroCluster IP configuration. The command sets the transition mode on the local and partner cluster.

Parameters

`-transition-mode <Mode For FC To Ip Transition>` - Transition Mode (privilege: advanced)

This parameter specifies the setting for the transition-mode.

`[-force {true|false}]` - Force enable (privilege: advanced)

This optional parameter forces the setting for the transition-mode.

Examples

The following example shows how to enable the non-disruptive transition mode:

```
clusA::> metrocluster transition enable transition-mode non-disruptive
clusA::> metrocluster transition show-mode
 Transition Mode
 --------------------------
 non-disruptive
```

The following example shows how to enable the disruptive transition mode:

```
clusA::> metrocluster transition enable transition-mode disruptive
clusA::> metrocluster transition show-mode
 Transition Mode
 --------------------------
disruptive
```

metrocluster transition show-mode

Display the transition mode

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

Description

The `metrocluster transition show-mode` command is used to display the current setting of transition.

Examples

The following example shows the output of the command when transition is not enabled:

```
clusA::> metrocluster transition show-mode
 Transition Mode
 --------------------------
 not-enabled
```

The following example shows the output of the command when non-disruptive transition is enabled:

```
clusA::> metrocluster transition show-mode
 Transition Mode
 --------------------------
 non-disruptive
```

The following example shows the output of the command when disruptive transition is enabled:

```
clusA::> metrocluster transition show-mode
 Transition Mode
 --------------------------
 disruptive
```

metrocluster vserver commands

metrocluster vserver recover-from-partial-switchback

Recover vservers from partial switchback

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

Description

The `metrocluster vserver recover-from-partial-switchback` command executes the necessary steps needed for a Vserver to be in healthy state after partial completion of the Switchback.

Parameters

Examples

```
cluster::> metrocluster vserver recover-from-partial-switchback
```

metrocluster vserver recover-from-partial-switchover

Recover vservers from partial switchover

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

Description

The `metrocluster vserver recover-from-partial-switchover` command executes the necessary steps needed for a Vserver to be in healthy state after partial completion of the Switchover.

Parameters

Examples

```
cluster::> metrocluster vserver recover-from-partial-switchover
```

metrocluster vserver resync

Resynchronize Vserver with its partner Vserver

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

Description

The `metrocluster vserver resync` command resynchronizes the Vserver with its partner Vserver

Parameters

`-cluster <Cluster name>` - Cluster Name

Name of the cluster where the Vserver belongs

`-vserver <vserver>` - Vserver

Name of the Vserver to be resynchronized

Examples

```
cluster::> metrocluster vserver resync -cluster clus1 -vserver vs1
```

metrocluster vserver show

Display MetroCluster Vserver relationships

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

Description

The `metrocluster vserver show` command displays configuration information for all pairs of Vservers in MetroCluster.

Parameters

{ [-fields <fieldname>,…]

The command output includes the specified field or fields

| [-creation-time] (privilege: advanced)

Shows the last configuration modification time on the Vserver

| [-instance] }

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

[-cluster <Cluster name>] - Cluster Name

Name of the cluster where the Vserver belongs

[-vserver <vserver>] - Vserver

Name of the Vserver

[-partner-vserver <vserver>] - Partner Vserver

Name of the partner Vserver

[-configuration-state {healthy|unhealthy|degraded|pending-setup|syncing|replication-paused|pending-switchback}] - Configuration State

Configuration states include:

- *healthy*
- *unhealthy*
- *degraded* indicates that Vservers are not in sync
- *syncing* indicates that the Vserver configuration is being synchronized
- *replication-paused* indicates that the configuration replication was manually paused
- *pending-setup* indicates that partner Vserver creation is pending

[-corrective-action <text>] - Corrective Action

Corrective action which can be followed to successfully create the partner Vserver

[-creation-time-of-last-applied-change <MM/DD/YYYY HH:MM:SS>] - Creation Time on the Source

Last configuration modification time on the Vserver

[-failed-row <smdb-row-op>] - Failed Row

Failed replication row

[-failed-object <text>] - Failed Object

Displays the failed object name

[-failed-reason <text>] - Failed Replication Reason

Failed replication reason

[-out-of-sync <true>] - Is Out of Sync

Indicates that the Vserver configuration replication is not in sync

[-config-resume-time <MM/DD/YYYY HH:MM:SS>] - Configuration Resume Time

Displays the resume time of the Vserver configuration replication

[-irrecoverable-apply-failed-reason <text>] - Reason for Failure

Reason for the replication failure

Examples

The following example shows the output of the command when partner Vservers are created

```
clusA::> metrocluster vserver show
Cluster: clusA
                                          Partner
Configuration
            Vserver                       Vserver
State
            -------------------           ----------------------
-----------------
            clusA                         clusB
healthy
            vs1                           vs1-mc
healthy
Cluster: clusB
                                          Partner
Configuration
            Vserver                       Vserver
State
            -------------------           ----------------------
-----------------
            clusB                         clusA
healthy
            3 entries were displayed.
```

The following example shows the output of the command when the partner Vserver creation is pending

```
clusA::> metrocluster vserver show
Cluster: clusA
                                          Partner
Configuration
            Vserver                       Vserver
State
            -------------------           ----------------------
-----------------
            clusA                         clusB
healthy
            vs1                           -
pending-setup
            Corrective Action: Create Ipspace ips1 on the partner cluster.
        2 entries were displayed.
```

network commands

network ping

Ping

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

Description

The `network ping` command displays whether a remote address is reachable and responsive, the (if specified) number of transmitted and received packets, and their round-trip time. The command requires a source node or logical interface from where the ping will be run, and a destination IP address. You can specify the source node by name, or a logical interface and its Vserver.

Parameters

{ -node <nodename> - Node

Use this parameter to send the ping from the node you specify.

| -lif <lif-name> - Logical Interface }

Use this parameter to send the ping from the logical interface you specify.

-vserver <vserver> - Vserver

Use this parameter to send the ping from the Vserver where the intended logical interface resides. The default value is the system Vserver for cluster administrators.

[-use-source-port {true|false}] - (DEPRECATED)-Use Source Port of Logical Interface (privilege: advanced)

This parameter is only applicable when the `-lif` parameter is specified. When set to true, the ping packet will be sent out via the port which is currently hosting the IP address of the logical interface. Otherwise, the ping packet will be sent out via a port based on the routing table.

> The `use-source-port` parameter is deprecated and may be removed in a future release of ONTAP.

-destination <Remote InetAddress> - Destination

Use this parameter to specify the remote internet address destination of the ping.

[-s, -show-detail <true>] - Show Detail Output

Use this parameter to display detailed output about the ping.

[-R, -record-route <true>] - Record Route

Use this parameter to display the route followed by the ping. You should set this option to `false` for pinging to succeed.

[-v, -verbose <true>] - Show All ICMP Packets

Use this parameter to display all ICMP packets.

[-packet-size <integer>] - Packet Size

Use this parameter to specify the number of data bytes to be sent in the ping packet. The default is 56 bytes, which is 64 ICMP data bytes total after 8 bytes of ICMP header data is added.

[-count <integer>] - Count

Use this parameter to specify the maximum number of ECHO_REQUESTS to be sent to the destination. The default is 20 requests. In the absence of the 'show-detail' option, ping reports that the destination is alive after receiving the first ECHO_REPLY response, independent of this value.

[-wait <integer>] - Packet Send Wait Time (secs)

Use this parameter to specify the number of seconds to wait between sending packets. The default is one second.

[-flood <true>] - Flood Ping (privilege: advanced)

Use this parameter to execute the command in flood mode. In flood mode, the command issues pings as fast as they are received, unless you specify a wait time.

[-D, -disallow-fragmentation <true>] - Disallow Packet Fragmentation

Use this parameter to prevent transport mechanisms from fragmenting ping packets in transit. Preventing fragmentation assures consistent packet size, making it easier to see transport bottlenecks.

[-wait-response <integer>] - Packet Response Wait Time (ms)

Use this parameter to specify the number of milliseconds to wait for each response packet. The default is 10000 milliseconds (10 seconds).

[-services <LIF Service Name>,…] - Services

Use this parameter to specify the list of services used to select the LIF from which to send the ping.

Examples

This example shows a ping from node xena to the destination server 10.98.16.164 with the server responding that it is up and running.

```
cluster1::> network ping -node xena -destination 10.98.16.164
(network ping)
10.98.16.164 is alive
```

network ping6

Ping an IPv6 address

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

Description

The `network ping6` command uses the ICMPv6 protocol's mandatory ICMP6_ECHO_REQUEST datagram to elicit an ICMP6_ECHO_REPLY from a host or gateway. ICMP6_ECHO_REQUEST datagrams ("pings") have an IPv6 header, and ICMPv6 header formatted as documented in RFC2463.

Parameters

{ -node <nodename> - Node Name

Use this parameter to originate ping6 from the specified node.

| -lif <lif-name> - Logical Interface }

Use this parameter to originate ping6 from the specified logical interface.

-vserver <vserver name> - Vserver Name

Use this parameter to originate ping6 from the specified Vserver. The default value is the system Vserver for cluster administrators.

-destination <Remote InetAddress> - Destination

Use this parameter to specify the IPv6 address of the destination node.

[-b, -buffer-size <integer>] - Socket Buffer Size

Use this parameter to set the socket buffer size.

[-c, -count <integer>] - Max Requests to Send/Recieve

Use this parameter to specify the maximum number of requests and replies. The default value is 20.

[-H, -reverse-lookup <true>] - Reverse-lookup of IPv6 addresses

Use this parameter to specify reverse-lookup of IPv6 addresses. Unless this parameter is specified, ping6 command does not attempt reverse lookup.

[-i, -interval <integer>] - Wait between Packets (secs)

Use this parameter to specify the delay time between packets in seconds. The default value is 1 second. This parameter is incompatible with the flood parameter.

[-l, -preload <integer>] - Send Packets as Fast as Possible (privilege: advanced)

Use this parameter if preload is required. If specified, ping6 sends that many packets as fast as possible before falling into its normal mode of behavior.

[-use-source-port {true|false}] - Use Source Port of Logical Interface (privilege: advanced)

This parameter is only applicable when the `-lif` parameter is specified. When set to true, the ping packet will be sent out via the port which is currently hosting the IP address of the logical interface. Otherwise, the ping packet will be sent out via a port based on the routing table.

[-p, -pattern <text>] - Up to 16 'pad' Specified for Out Packet

Use this parameter to fill the -16 'pad' bytes in the sent packet. This is useful for diagnosing data dependent problems in a network. For example, -pattern ff causes the sent packet to be filled with all ones.

[-packet-size <integer>] - Packet Size

Use this parameter to specify the number of data bytes to be sent. The default is 56, which translates to 64 ICMP data bytes when combined with the 8 bytes of ICMP header data.

[-v, -verbose <true>] - Show All ICMP Packets

Use this parameter to get verbose output. Verbose output displays both source address and destination addresses. Received ICMP packets other than ECHO_RESPONSE are listed. This parameter can be used only in conjunction with the show-detail parameter.

[-s, -show-detail <true>] - Show Detail Output

Use this parameter to display detailed output about the ping.

[-f, -flood <true>] - Flood Ping (privilege: advanced)

Use this parameter to output packets as fast as they come back or one hundred times per second, whichever is more. For every ECHO_REQUEST sent a period "." is printed, while for every ECHO_REPLY received a backspace is printed. This provides a rapid display of how many packets are being dropped. This can be very hard on a network and should be used with caution.

[-D, -disable-fragmentation <true>] - Disable Fragmentation.

Use this parameter to disallow fragmentation of the outgoing packets, if they do not fit in the Maximum Transmission Unit.

Examples

This example shows a ping6 from node 'node1' to the destination server ipv6.google.com with the server responding that it is up and running.

```
cluster1::> network ping6 -node node1 -destination ipv6.google.com
    ipv6.google.com is alive.
```

network test-path

Test path performance between two nodes

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

Description

The `network test-path` command runs a performance test between two nodes. The command requires a source node, destination node, destination cluster, and application, or session type. All tests are run using intracluster or intercluster LIFs, depending on whether the test is between two nodes in the same cluster, or between nodes in peered clusters.

The test itself is different from most bandwidth test tools. It creates a "session" consisting of TCP connections between all possible paths between the nodes being tested. This is how internal ONTAP applications communicate between nodes. This means the test is using multiple paths, and thus the bandwidth reported might exceed the capacity of a single 10 Gb path.

Parameters

-source-node {<nodename>|local} - Node Initiating Session (privilege: advanced)

Use this parameter to specify the node that initiates the test. Source-node parameter must be a member of the cluster in which the command is run.

-destination-cluster <Cluster name> - Cluster Containing Passive Node (privilege: advanced)

Use this parameter to specify the destination cluster; the local cluster, or a peered cluster.

`-destination-node <text>` - Remote Node in Destination Cluster (privilege: advanced)

Use this parameter to specify the destination node in the destination cluster

`-session-type {AsyncMirrorLocal|AsyncMirrorRemote|RemoteDataTransfer}` - Type of Session to Test (privilege: advanced)

The session type parameter is used to mimic the application settings used. A session consists of multiple TCP connections.

- AsyncMirrorLocal: settings used by SnapMirror between nodes in the same cluster
- AsyncMirrorRemote: settings used by SnapMirror between nodes in different clusters
- RemoteDataTransfer: settings used by ONTAP for remote data access between nodes in the same cluster

The default session-type is AsyncMirrorRemote.

`[-connection-detail {true|false}]` - Display Test Results per Network Path (privilege: advanced)

Boolean argument to specify if the network test-path run should be executed for specific network paths. This defaults to false.

Examples

The following example runs a test between two nodes in the same cluster:

```
cluster1::*> network test-path -source-node node1 -destination-cluster
cluster1 -destination-node node2
Test Duration: 10.65 secs
   Send Throughput: 1092.65 MB/sec
Receive Throughput: 1092.65 MB/sec
           MB Sent: 11633.69
       MB Received: 11633.69
       Avg Latency:    64.40 ms
       Min Latency:     2.41 ms
       Max Latency:  2099.17 ms
```

network traceroute

Traceroute

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

Description

The `network traceroute` command performs a network probe from a node to a specified IP address. The command requires a source node or logical interface and a destination IP address. You can specify the source node by name, or specify a logical interface and its Vserver. The traceroute is performed between the source and destination.

Parameters

{ -node <nodename> - Node

Use this parameter to originate the traceroute from the node you specify.

| -lif <lif-name> - Logical Interface }

Use this parameter to originate the traceroute from the specified network interface.

-vserver <vserver> - LIF Owner

Use this parameter to originate the traceroute from the Vserver where the intended logical interface resides. The default value is the system Vserver for cluster administrators.

-destination <Remote InetAddress> - Destination

Use this parameter to specify the remote internet address destination of the traceroute.

[-m, -maxttl <integer>] - Maximum Number of Hops

Use this parameter to specify the maximum number of hops (time-to-live) setting used by outgoing probe packets. The default is 30 hops.

[-n, -numeric <true>] - Print Hop Numerically

Use this parameter to print the hop addresses only numerically rather than symbolically and numerically.

[-port <integer>] - Base UDP Port Number

Use this parameter to specify the base UDP port number used in probes. The default is port 33434.

[-packet-size <integer>] - Packet Size

Use this parameter to specify the size of probe packets, in bytes.

[-q, -nqueries <integer>] - Number of Queries

Use this parameter to specify the number of probes per hop. The default is 3 probes.

[-v, -verbose <true>] - Verbose Output

Use this parameter to display all received ICMP packets, rather than just TIME_EXCEEDED and UNREACHABLE packets.

[-w, -waittime <integer>] - Wait Between Packets (secs)

Use this parameter to specify the time (in seconds) to wait for the response to a probe. The default is 5 seconds.

Examples

This example shows a traceroute from node node1 to a destination address of 10.98.16.164, showing a maximum of five hops.

```
cluster1::> traceroute -node node1 -destination 10.98.16.164 -maxttl 5
 1   10.68.208.1 <10.68.208.1> 0.307 ms 293 ms 305 ms
 2   152.164.13.205 <152.164.13.205> 3.754 ms 3.722 ms 3.981 ms
 3   68.137.122.222 <68.137.122.222> 25.603 ms 24.947 ms 24,565 ms
 4   * * *
 5   * * *

traceroute to 10.98.16.164, 5 hops max, 52 byte packets
```

network traceroute6

traceroute6

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

Description

The `network traceroute6` command performs a network probe from a node to a specified IPv6 address. The command requires a source node or logical interface, Vserver from where traceroute6 will originate and a destination IPv6 address. traceroute6 is performed between the source and destination.

Parameters

{ -node <nodename> - Node

Use this parameter to originate traceroute6 from the node you specify. This parameter is available only to cluster administrators.

| -lif <lif-name> - Logical Interface }

Use this parameter to originate traceroute6 from the logical interface you specify.

-vserver <vserver name> - LIF Owner

Use this parameter to originate traceroute6 from the Vserver you specify. The default value is the system Vserver for cluster administrators.

[-d, -debug-mode <true>] - Debug Mode

Use this parameter to enable socket level debugging. The default value is false.

{ [-I, -icmp6 <true>] - ICMP6 ECHO instead of UDP

Use this parameter to specify the use of ICMP6 ECHO instead of UDP datagrams for the probes. The default value is false.

| [-U, -udp <true>] - UDP }

Use this parameter to specify the use of UDP datagrams for the probes. The default value is true.

[-n, -numeric <true>] - Print Hops Numerically

Use this parameter to print the hop addresses only numerically rather than symbolically and numerically. The default value is false.

[-v, -verbose <true>] - Verbose Output

Use this parameter to display all received ICMP packets, rather than just TIME_EXCEEDED and UNREACHABLE packets. The default value is false.

[-f, -first-hop <integer>] - Number of Hops to Skip in Trace

Use this parameter to specify the number of hops to skip in trace. The default value is 1.

[-g, -gateway <Remote InetAddress>] - Intermediate Gateway

Use this parameter to specify the intermediate gateway.

[-m, -hop-limit <integer>] - Maximum Number of Hops

Use this parameter to specify the maximum hoplimit, up to 255. The default value is 64 hops.

[-p, -port <integer>] - Base UDP Port Number

Use this parameter to specify the base UDP port number used in probes. The default value is port 33434.

[-q, -nqueries <integer>] - Number of Queries

Use this parameter to specify the number of probes per hop. The default value is 3 probes.

[-w, -wait-time <integer>] - Wait Between Packets (secs)

Use this parameter to specify the delay time between probes in seconds. The default value is 5 seconds.

-destination <Remote InetAddress> - Destination

Use this parameter to specify the remote IPv6 address destination of traceroute6.

[-packet-size <integer>] - Packet Size

Use this parameter to specify the size of probe packets, in bytes. The default value is 16 bytes for ICMP6 ECHO and 12 bytes for UDP datagrams.

Examples

The following example shows traceroute6 from node node1 to the destination fd20:8b1e:b255:4071:d255:1fcd:a8cd:b9e8.

```
cluster1::> network traceroute6 -node node1 -vserver vs1
                -destination 3ffe:b00:c18:1::10
traceroute6 to 3ffe:b00:c18:1::10 (3ffe:b00:c18:1::10)
              from 2001:0db8:0000:f101::2,
           64 hops max, 12 byte packets
1  2001:0db8:0000:f101::1 4.249 ms  2.021 ms  0.864 ms
2  3ffe:2000:0:400::1  0.831 ms  0.579 ms
3  3ffe:2000:0:1::132  227.693 ms  227.596 ms  227.439 ms
4  3ffe:c00:8023:2b::2 229.028 ms  228.267 ms  231.891 ms
5  3ffe:2e00:e:c::3  227.929 ms  228.696 ms  228.558 ms
6  3ffe:b00:c18:1::10  227.702 ms 227.806 ms  227.439 ms
```

network arp commands

network arp create

Create static ARP entry

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

Description

The `network arp create` command creates a static ARP entry for a given Vserver. Statically created ARP entries will be stored permanently in the Vserver context and will be used by the network stack.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the name of the Vserver on which the ARP entry is created.

`-remotehost <IP Address>` - Remote IP Address

Use this parameter to specify the IP address to be added as an ARP entry.

`-mac <MAC Address>` - MAC Address

Use this parameter to specify the MAC address (Ethernet address) for the host specified with -remotehost. Specify the MAC address as six hex bytes separated by colons.

Examples

The following example creates a static ARP entry on Vserver vs1 for the remote host with the IP address 10.63.0.2 having MAC address 40:55:39:25:27:c1

```
cluster1::> network arp create -vserver vs1 -remotehost 10.63.0.2 -mac
40:55:39:25:27:c1
```

network arp delete

Delete static ARP entry

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

Description

The `network arp delete` command deletes static ARP entries from the Vserver and from the network stack.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the name of the Vserver from which the ARP entry is deleted.

`-remotehost <IP Address>` - Remote IP Address

Use this parameter to specify the IP address of the ARP entry being deleted.

Examples

The following example deletes the ARP entry for IP address 10.63.0.2 from the Vserver vs1.

```
cluster1::> network arp delete -vserver vs1 -remotehost 10.63.0.2
```

network arp show

Display static ARP entries

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

Description

The `network arp show` command displays static ARP entries present in a given Vserver. This command will not display dynamically learned ARP entries in the network stack. Use the network arp active-entry show command to display dynamically learned ARP entries in the network stack.

Parameters

{ [-fields <fieldname>,…]

Use this parameter to display only certain fields of the ARP table.

| [-instance] }

Use this parameter to display all the fields of the ARP table.

[-vserver <vserver name>] - Vserver Name

Use this parameter to display ARP entries that are specific to a given Vserver.

[-remotehost <IP Address>] - Remote IP Address

Use this parameter to display ARP entries for the specified IP address

[-mac <MAC Address>] - MAC Address

Use this parameter to display ARP entry for the specified MAC address

[-ipspace <IPspace>] - IPspace

Use this parameter to specify the IPspace associated with the Vserver

Examples

The following example displays static ARP entries from the Vserver vs1.

```
cluster1::> network arp show -vserver vs1
Vserver     Remote Host       MAC Address
----------- ----------------- -----------------
vs1
            10.238.0.2        40:55:39:25:27:c1
```

Related Links

- network arp active-entry show

network arp active-entry delete

Delete active ARP entry from a System or Admin Vserver

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

Description

The `network arp active-entry delete` command deletes dynamically learned ARP entries from the network stack of a node. To delete statically configured ARP entries use the network arp delete command.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the name of the node in which the ARP entry is deleted.

`-vserver <vserver>` - System or Admin Vserver Name (privilege: advanced)

Use this parameter to specify the name of the Vserver in which the ARP entry is deleted. Only Vservers with a type of Admin or System have dynamically learned ARP entries.

`-subnet-group <IP Address/Mask>` - Subnet Group Name (privilege: advanced)

Use this parameter to specify the name of the routing group in which the ARP entry is deleted.

`-remotehost <text>` - Remote IP Address (privilege: advanced)

Use this parameter to specify the IP address to be deleted from the active ARP entries.

`-port <text>` - Port (privilege: advanced)

Use this parameter to specify the name of the Port to be deleted from the active ARP entries.

Examples

The following example deletes the active ARP entry with an IP address of 10.224.64.1, subnet group of 0.0.0.0/0, port e0c on node node2 in the Admin Vserver cluster1:

```
cluster1::*> network arp active-entry delete -node cluster1-01 -vserver
cluster1 -subnet-group 0.0.0.0/0 -remotehost 10.224.64.1 -port e0c
```

Related Links

- network arp delete

network arp active-entry show

Display active ARP entries organized by Vserver

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

Description

The `network arp active-entry show ` command displays ARP entries present in the network stack of the node. The entries includes both dynamically learned ARP entries and user configured static ARP entries.

Parameters

`[-fields <fieldname>,…]`

Use this parameter to display only certain fields of the active ARP table.

| `[-instance ]` (privilege: advanced) }

Use this parameter to display all the fields of the active ARP table.

`[-node {<nodename>|local}]` - Node (privilege: advanced)

Use this parameter to display active ARP entries that are specific to a given node.

`[-vserver <vserver>]` - System or Admin Vserver Name (privilege: advanced)

Use this parameter to display active ARP entries that are specific to a given System or Admin Vserver. Data and Node Vservers will not have any active-arp entries.

`[-subnet-group <IP Address/Mask>]` - Subnet Group Name (privilege: advanced)

Use this parameter to display active ARP entries that are specific to a given subnet group.

`[-remotehost <text>]` - Remote IP Address (privilege: advanced)

Use this parameter to display active ARP entries for the specified IP address.

`[-port <text>]` - Port (privilege: advanced)

Use this parameter to display active ARP entries for the specified Port name.

`[-mac <text>]` - MAC Address (privilege: advanced)

Use this parameter to display the active ARP entry for the specified MAC address.

`[-ipspace <IPspace>]` - IPspace (privilege: advanced)

Use this parameter to specify the IPspace associated with the System or Admin Vserver.

Examples

The following example displays active ARP entries for the Admin Vserver cluster1:

```
cluster1::*> network arp active-entry show -vserver cluster1

Node: node-01
Vserver: cluster1
Subnet Group: 169.254.0.0/16
Remote IP Address MAC Address       Port
----------------- ----------------- -------
169.254.106.95    0:55:39:27:d1:c1  lo
```

network bgp commands

network bgp config create

Create BGP configuration

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

Description

The `network bgp config create` command is used to create the border gateway protocol (BGP) configuration for a node. It can be used to override the BGP parameters defined in the global BGP defaults.

Parameters

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

This parameter specifies the node on which configuration details will be created.

`-asn <integer>` - Autonomous System Number

This parameter specifies the autonomous system number (ASN). The ASN attribute is a positive integer of the range from 1 to 4,294,967,295. It should typically be chosen from RFC6996 "Autonomous System (AS) Reservation for Private Use" or the AS number assigned to the operator’s organization.

`-hold-time <integer>` - Hold Time

This parameter specifies the hold time in seconds. The default value is 180.

`-router-id <IP Address>` - Router ID

This parameter specifies the local router ID. The router-id value takes the form of an IPv4 address. The default router-id will be initialized using a local IPv4 address in admin vserver.

Examples

```
cluster1::> network bgp config create -node node1 -asn 10 -hold-time 180
-router-id 10.0.1.112
```

network bgp config delete

Delete BGP configuration

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

Description

The `network bgp config delete` command deletes a node's border gateway protocol (BGP) configuration. A BGP configuration cannot be deleted if there are BGP peer groups configured on the associated node.

Parameters

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

This parameter specifies the node for which the BGP configuration will be deleted.

Examples

```
cluster1::> network bgp config delete -node node1
```

network bgp config modify

Modify BGP configuration

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

Description

The `network bgp config modify` command is used to modify a node's border gateway protocol (BGP) configuration.

Parameters

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

This parameter specifies the node on which BGP configuration will be modified.

`[-asn <integer>]` - Autonomous System Number

This parameter specifies the autonomous system number (ASN). The ASN attribute is a positive integer of the range from 1 to 4,294,967,295. It should typically be chosen from RFC6996 "Autonomous System (AS) Reservation for Private Use" or the AS number assigned to the operator's organization.

`[-hold-time <integer>]` - Hold Time

This parameter specifies the hold time in seconds.

`[-router-id <IP Address>]` - Router ID

This parameter specifies the local router ID. The router-id value takes the form of an IPv4 address.

Examples

```
cluster1::> network bgp config modify -node node1 -router-id 1.1.1.1 -asn
20
```

network bgp config show

Display BGP configuration

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

Description

The `network bgp config show` command displays the border gateway protocol (BGP) configuration for each node.

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>|local}] - Node

This parameter selects the BGP configurations that match the specified node.

[-asn <integer>] - Autonomous System Number

This parameter selects the BGP configurations that match the specified autonomous system number.

[-hold-time <integer>] - Hold Time

This parameter selects BGP configurations that match the specified hold time.

[-router-id <IP Address>] - Router ID

This parameter selects the BGP configurations that match the specified router ID.

Examples

```
cluster1::> network bgp config show
        Autonomous
        System    Hold Time
    Node    Number    (seconds)  Router ID
    ---------------  ---------------  ---------------  ---------
    node1    10    180    10.0.1.112
```

network bgp defaults modify

Modify BGP defaults

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

Description

The `network bgp defaults modify` command modifies the global defaults for border gateway protocol (BGP) configurations.

Parameters

`[-asn <integer>]` - Autonomous System Number

This parameter specifies the autonomous system number (ASN). The ASN attribute is a positive integer. It should typically be chosen from RFC6996 "Autonomous System (AS) Reservation for Private Use", or the AS number assigned to the operator's organization. The default ASN is 65501.

`[-hold-time <integer>]` - Hold Time

This parameter specifies the hold time in seconds. The default value is 180.

Examples

```
cluster1::> network bgp defaults modify -asn 20
```

network bgp defaults show

Display BGP defaults

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

Description

The `network bgp defaults show` command displays the global defaults for border gateway protocol (BGP) configurations.

Examples

```
cluster1::> network bgp defaults show
Autonomous
System Number  Hold Time
    (Seconds)
-------------  --------------
10    180
```

network bgp peer-group create

Create a new BGP peer group

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

Description

The `network bgp peer-group create` command is used to create a border gateway protocol (BGP) peer group. A BGP peer group will advertise VIP routes for the list of vservers in the peer group's `vserver-list` using the BGP LIF of the peer group. A BGP peer group will advertise VIP routes to a peer router using the border gateway protocol. The address of the peer router is identified by the `peer-address` value.

Parameters

`-ipspace <IPspace>` - IPspace Name

This parameter specifies the IPspace of the peer group being created.

`-peer-group <text>` - Peer Group Name

This parameter specifies the name of the peer group being created.

`-bgp-lif <lif-name>` - BGP LIF

This parameter specifies the BGP interface (BGP LIF) of the peer group being created.

`-peer-address <IP Address>` - Peer Router Address

This parameter specifies the IP address of the peer router for the peer group being created.

`[-peer-asn <integer>]` - Peer Router Autonomous number

This parameter specifies the peer router autonomous system number (ASN) in the peer group being created. The default value is the value of the local node's ASN.

`-route-preference <integer>` - Route Preference

This parameter specifies the preference field in BGP update messages for VIP routes. If a router receives multiple VIP route announcements for the same VIP LIF from different BGP LIFs, it will install the one that has the highest preference value. The default route preference value is 100.

`[-asn-prepend-type <ASN Prepend type>]` - ASN prepend type (privilege: advanced)

This parameter specifies the ASN that will be prepended in the BGP attributes. The possible values are `local-asn` and `peer-asn` . The default behavior is not to prepend any ASN.

`[-asn-prepend-count <integer>]` - ASN prepend count (privilege: advanced)

This parameter specifies the number of times ASN, as specified in `asn-prepend-type` will be prepended in the BGP path attributes. The default behavior is not to prepend any ASN.

`[-community <BGP community>,…]` - BGP Community (privilege: advanced)

This parameter specifies the communities that will be included in the BGP path attributes. The default behavior is not to include any community in BGP path attributes.

`[-med <integer>]` - Multi Exit Discriminator (privilege: advanced)

This parameter specifies the Multi Exit Discriminator (MED) attribute of BGP update messages, which can be used by routers for best path selection, in cases where more than one peer advertises the same route with similar attributes.

[-use-peer-as-next-hop {true|false}] - Use Peer Address As Next Hop

This parameter specifies whether the peer group uses the peer address as a next hop route. When the value is true, the peer address is used as the next hop router for packets sent from VIP LIFs via the port on which `bgp-lif` is configured. Internally, a default route with a gateway configured as the `peer-address` is added automatically on the node for all the Vservers in this peer group's IPspace. The route will be added for a Vserver only if it has a VIP LIF hosted on the current node of `bgp-lif` . Note that these automatically installed default routes are for VIP traffic; however, they can be used for non-VIP traffic as well if a Vserver hosts both VIP and non-VIP LIFs in the same subnet as `bgp-lif` . This route will have metric of 20 and will be used to forward traffic through the current port of `bgp-lif` . The default value of this parameter is false.

[-md5-enabled {true|false}] - TCP MD5 Enabled (privilege: advanced)

This parameter specifies whether the peer group configuration uses TCP MD5 authentication to communicate with its peer router. When the value is true, an `md5-secret` must be specified for the peer group being created. The default value is false.

[-md5-secret <text>] - TCP MD5 Secret Key (privilege: advanced)

This parameter specifies TCP MD5 secret key. The value for `md5-secret` may be specified as either a string of printable ASCII characters or as a hexadecimal string. If a string of ASCII characters is specified, the string should be enclosed in quotes and be from 1-79 characters in length. If a hexadecimal string is specified, the string of hexadecimal bytes should be prefaced with 0x (e.g., 0x893785c6) and contain from 2-160 hexadecimal characters. The `md5-secret` should match the MD5 secret configured on the peer router.

Examples

```
cluster1::> network bgp peer-group create -peer-group group1 -ipspace
Default -bgp-lif bgp_lif -peer-address 10.0.1.112
```

network bgp peer-group delete

Delete a BGP peer group

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

Description

The `network bgp peer-group delete` command is used to delete border gateway protocol (BGP) peer group configuration.

Parameters

-ipspace <IPspace> - IPspace Name

This parameter specifies the IPspace of the BGP peer group being deleted.

-peer-group <text> - Peer Group Name

This parameter specifies the name of the BGP peer group being deleted.

Examples

```
cluster1::> network bgp peer-group delete -ipspace Default -peer-group
group1
```

network bgp peer-group modify

Modify a BGP peer group

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

Description

The `network bgp peer-group modify` command is used to modify a border gateway protocol (BGP) peer group configuration.

Parameters

`-ipspace <IPspace>` - IPspace Name

This parameter specifies the IPspace of the peer group being modified.

`-peer-group <text>` - Peer Group Name

This parameter specifies the name of the peer group being modified.

`[-peer-address <IP Address>]` - Peer Router Address

This parameter specifies an updated value for the IP address of the peer router.

`[-use-peer-as-next-hop {true|false}]` - Use Peer Address As Next Hop

This parameter specifies whether the peer group uses the peer address as a next hop route. When the value is true, the peer address is used as the next hop router for packets sent from VIP LIFs via the port on which `bgp-lif` is configured. Internally, a default route with a gateway configured as the `peer-address` is added automatically on the node for all the Vservers in this peer group's IPspace. The route will be added for a Vserver only if it has a VIP LIF hosted on the current node of `bgp-lif` . Note that these automatically installed default routes are for VIP traffic; however, they can be used for non-VIP traffic as well if a Vserver hosts both VIP and non-VIP LIFs in the same subnet as `bgp-lif` . This route will have metric of 20 and will be used to forward traffic through the current port of `bgp-lif` . The default value of this parameter is false.

`[-md5-enabled {true|false}]` - TCP MD5 Enabled (privilege: advanced)

This parameter specifies whether the peer group configuration uses TCP MD5 authentication to communicate with its peer router. When the value is true, an `md5-secret` must be specified for the peer group being created. The default value is false.

`[-md5-secret <text>]` - TCP MD5 Secret Key (privilege: advanced)

This parameter specifies TCP MD5 secret key. The value for `md5-secret` may be specified as either a string of printable ASCII characters or as a hexadecimal string. If a string of ASCII characters is specified, the string should be enclosed in quotes and be from 1-79 characters in length. If a hexadecimal string is specified, the string of hexadecimal bytes should be prefaced with 0x (e.g., 0x893785c6) and contain from 2-160 hexadecimal characters. The `md5-secret` should match the MD5 secret configured on the peer router.

Examples

```
cluster1::> network bgp peer-group modify -ipspace Default -peer-group
peer1 -peer-address 10.10.10.10
```

network bgp peer-group rename

Rename a BGP peer group

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

Description

The `network bgp peer-group rename` command is used to assign a new name to a BGP peer group.

Parameters

`-ipspace <IPspace>` - IPspace Name (privilege: advanced)

This parameter specifies the IPspace of the peer group being renamed.

`-peer-group <text>` - Peer Group Name (privilege: advanced)

The name of the peer group to be updated.

`-new-name <text>` - New Name (privilege: advanced)

The new name for the peer group.

Examples

```
cluster1::> network bgp peer-group rename -peer-group old_name -new-name
new_name
```

network bgp peer-group show

Display BGP peer groups information

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

Description

The `network bgp peer-group show` command displays the BGP peer groups configuration.

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.

[-ipspace <IPspace>] - IPspace Name

This parameter selects peer groups that match the specified IPspace.

[-peer-group <text>] - Peer Group Name

This parameter selects peer groups that match the specified name.

[-bgp-lif <lif-name>] - BGP LIF

This parameter selects peer groups that match the specified BGP Interface.

[-peer-address <IP Address>] - Peer Router Address

This parameter selects peer groups that match the specified peer router address.

[-peer-asn <integer>] - Peer Router Autonomous number

This parameter selects peer groups that match the specified autonomous system number.

[-state <BGP Session State>] - Peer Group State

This parameter selects peer groups that match the specified BGP session state.

[-bgp-node <nodename>] - BGP LIF Node

This parameter selects peer groups that match the specified `bgp-node` value. This value is calculated based on the current node of the corresponding BGP LIF.

[-bgp-port <netport>] - BGP LIF Port

This parameter selects peer groups that match the specified `bgp-port` value. This value is calculated based on the current port of the associated BGP LIF.

[-route-preference <integer>] - Route Preference

This parameter selects peer groups that match the specified route preference value.

[-asn-prepend-type <ASN Prepend type>] - ASN prepend type (privilege: advanced)

This parameter selects peer groups that match the specified `asn-prepend-type` value. The possible values are `local-asn` and `peer-asn` .

[-asn-prepend-count <integer>] - ASN prepend count (privilege: advanced)

This parameter selects peer groups that match the specified `asn-prepend-count` value.

[-community <BGP community>,…] - BGP Community (privilege: advanced)

This parameter selects peer groups that match the specified `community` value.

[-med <integer>] - Multi Exit Discriminator (privilege: advanced)

This parameter selects peer groups that match the specified med value.

[-use-peer-as-next-hop {true|false}] - Use Peer Address As Next Hop

This parameter selects peer groups that match the specified use-peer-as-next-hop value.

[-md5-enabled {true|false}] - TCP MD5 Enabled (privilege: advanced)

This parameter selects peer groups that match the specified md5-enabled value.

[-md5-secret <text>] - TCP MD5 Secret Key (privilege: advanced)

This parameter selects peer groups that match the specified TCP MD5 secret key value. The value for `md5-secret` may be specified as either a string of printable ASCII characters or as a hexadecimal string. If a string of ASCII characters is specified, the string should be enclosed in quotes and be from 1-79 characters in length. If a hex string is specified, the string of hexadecimal bytes should be prefaced with 0x (e.g., 0x893785c6) and be from 2-160 hexadecimal characters in length.

Examples

```
cluster1::> network bgp peer-group show
    IPspace: Default
    Peer        Local BGP Peer router                      Autonomous
    Group       Interface Address/subnet    state          Number     Node
Port
    ----------  --------- ----------------- -------------- ----------
------ -----
    gp1         bgp_lif1  10.0.5.37         up                     10
node1  e1a
    gp2         bgp_lif2  10.0.6.38         up                     12
node1  e2a
```

network bgp vserver-status show

Display Vserver BGP status

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

Description

The `network bgp vserver-status show` command displays the per-node border gateway protocol (BGP) status for each vserver. The BGP status for a particular vserver is "up" when at least one BGP peer group supporting that vserver is able to communicate with its peer router.

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>|local}] - Node

This parameter selects the BGP status that match the specified node.

[-vserver <vserver name>] - Vserver

This parameter selects the BGP status for specified vserver.

[-ipv4-status {unknown|unconfigured|up|down}] - IPv4 status

This parameter selects the BGP status that matches the specified status for IPv4 address family.

[-ipv6-status {unknown|unconfigured|up|down}] - IPv6 status

This parameter selects the BGP status that matches the specified status for IPv6 address family.

Examples

```
cluster1::> network bgp vserver-status show
    Node                vserver    IPv4 status IPv6 status
    ------------------- ---------- ----------- -----------
    node1               vs1        up          up
```

network cloud commands

network cloud routing-table create

Create a new external routing table

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

Description

The `network cloud routing-table create` command creates a new external routing table.

Parameters

-route-table-id <text> - Route Table ID (privilege: advanced)

This parameter is used to provide the name of the external routing table to be created.

Examples

The following example creates an external routing table "eni-123456":

```
cluster1::> network cloud routing-table create -route-table-id eni-123456
```

network cloud routing-table delete

Delete an existing external routing table

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

Description

The `network cloud routing-table delete` deletes an existing external routing table.

Parameters

`-route-table-id <text>` - Route Table ID (privilege: advanced)

This parameter is used to provide the name of an existing external routing table to be deleted.

Examples

The following example deletes the external routing table "eni-123456":

```
cluster1::> network cloud routing-table delete -route-table-id eni-123456
```

network cloud routing-table show

Show existing external routing tables

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

Description

The `network cloud routing-table show` command retrieves the configured routing tables on mediator and displays them.

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.

`[-route-table-id <text>]` - Route Table ID (privilege: advanced)

This parameter is the name of the external routing table to be shown.

Examples

The following example shows external routing tables, for this cluster, that are configured on the mediator

```
cluster1::> network cloud routing-table show
Route Table ID
 ----------------------
rtb-16924571
rtb-9c9245fb
rtb-a36ca1c4
3 entries were displayed.
```

network connections commands

network connections active show-clients

Show a count of the active connections by client

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

Description

The `network connections active show-clients` command displays information about client connections, including the client's IP address and the number of client connections.

> The results of this command set are refreshed independently every 30 seconds and might not reflect the immediate state of the system.

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>|local}] - Node

Use this parameter to display information only about the connections on the node you specify.

[-vserver <vserver>] - Vserver

This parameter is used by the system to break down the output per vserver.

[-remote-address <Remote IP>] - Remote IP Address

Use this parameter to display information only about the connections that use the remote IP address you specify.

[-count <integer>] - Client Count

Use this parameter to only clients with the number of active client connections you specify.

Examples

The following example displays information about active client connections:

```
cluster1::> network connections active show-clients
Node    Vserver Name    Client IP Address     Count
------  --------------  -----------------     ------
node0   vs1             192.0.2.253                1
        vs2             192.0.2.252                2
        vs3             192.0.2.251                5
node1   vs1             192.0.2.250                1
        vs2             192.0.2.252                3
        vs2             customer.example.com       4
```

network connections active show-lifs

Show a count of the active connections by logical interface

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

Description

The `network connections active show-lifs` command displays the number of active connections on each logical interface, organized by node and Vserver.

> The results of this command set are refreshed independently every 30 seconds and might not reflect the immediate state of the system.

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>|local}] - Node

Use this parameter to display information only about the connections on the node you specify.

[-vserver <vserver>] - Vserver

Use this parameter to display information only about the connections that are using the node or Vserver you specify.

[-lif-name <lif-name>] - Logical Interface Name

Use this parameter to display information only about the connections that are using the logical interface you specify.

[-count <integer>] - Client Count

Use this parameter to display only logical interfaces with the number of active client connections you specify.

[-blocked-count <integer>] - (DEPRECATED)-Load Balancing Blocking Count

> This parameter has been deprecated and may be removed in a future version of ONTAP.

Use this parameter to display information only about data logical interfaces blocked from migrating and the connection that is blocking it.

Examples

The following example displays information about the servers and logical interfaces being used by all active connections:

```
cluster1::> network connections active show-lifs
Node      Vserver Name  Interface Name  Count
--------  ------------  --------------- ------
node0
          vs0           datalif1             3
          vs0           cluslif1             6
          vs0           cluslif2             5
node1
          vs0           datalif2             3
          vs0           cluslif1             3
          vs0           cluslif2             5
node2
          vs1           datalif2             1
          vs1           cluslif1             5
          vs1           cluslif2             3
node3
          vs1           datalif1             1
          vs1           cluslif1             2
          vs1           cluslif2             1
```

At privilege levels above "admin", the command displays an extra column.

```
cluster1::*> network connections active show-lifs
                                                 LB Migrate
Node      Vserver Name  Interface Name  Count  Blocking
--------  ------------  --------------- ------ ----------
node0
          vs0           datalif1             3          0
          vs0           cluslif1             6          0
          vs0           cluslif2             5          2
node1
          vs0           datalif2             3          0
          vs0           cluslif1             3          0
          vs0           cluslif2             5          0
node2
          vs1           datalif2             1          0
          vs1           cluslif1             5          0
          vs1           cluslif2             3          2
node3
          vs1           datalif1             1          0
          vs1           cluslif1             2          0
          vs1           cluslif2             1          0
```

network connections active show-protocols

Show a count of the active connections by protocol

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

Description

The `network connections active show-protocols` command displays the number of active connections per protocol, organized by node.

> The results of this command set are refreshed independently every 30 seconds and might not reflect the immediate state of the system.

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>|local}] - Node

Use this parameter to display information only about the connections on the node you specify.

[-vserver <vserver>] - Vserver

This parameter is used by the system to break down the output per vserver.

[-proto {UDP|TCP|TLS}] - Protocol

Use this parameter to display information only about the connections that use the network protocol you specify. Possible values include tcp (TCP), udp (UDP), and NA (not applicable).

[-count <integer>] - Client Count

Use this parameter to display only protocols with the number of active client connections you specify.

Examples

The following example displays information about all network protocols being used by active connections:

```
cluster1::> network connections active show-protocols
Node      Vserver Name    Protocol   Count
-------  --------------  ---------  ------
node0
         vs1             UDP            19
         vs1             TCP            11
         vs2             UDP            17
node1
         vs1             UDP            14
         vs2             TCP            10
```

network connections active show-services

Show a count of the active connections by service

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

Description

The `network connections active show-services` command displays the number of active connections by protocol service, organized by node.

> The results of this command set are refreshed independently every 30 seconds and might not reflect the immediate state of the system.

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>|local}] - Node

Use this parameter to display information only about the connections on the node you specify.

[-vserver <vserver>] - Vserver

This parameter is used by the system to break down the output per vserver

[-service <protocol service>] - Protocol Service

Use this parameter to display information only about the connections that use the protocol service you specify. Possible values include: nfs, iscsi, and loopback.

[-count <integer>] - Client Count

Use this parameter to display information only about protocol services with the number of active client connections you specify.

Examples

The following example displays information about all protocol services being used by active connections:

```
cluster1::> network connections active show-services
Node       Vserver Name    Service    Count
---------  --------------  ---------  ------
node0
           vs1             mount           3
           vs1             nfs            14
           vs1             nlm_v4          4
           vs1             cifs_srv        3
           vs1             port_map       18
           vs2             rclopcp        27
node1
           vs1             nfs             5
           vs2             rclopcp        12
           vs2             nfs             4
           vs2             port_map        8
```

network connections active show

Show the active connections in this cluster

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

Description

The `network connections active show` command displays information about active network connections.

> The results of this command set are refreshed independently every 30 seconds and might not reflect the immediate state of the system.

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.

| [-print-ip-addresses]

Print IP addresses for remote hosts — do not attempt to resolve the addresses to a hostname.

| [-instance] }

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

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

Selects the connections that match this parameter value.

[-cid <Cid>] - Connection ID

Selects the connections that match this parameter value.

[-vserver <vserver>] - Vserver

Selects the connections that match this parameter value.

[-lif-name <lif-name>] - Logical Interface Name

Selects the connections that match this parameter value.

[-local-address <IP Address>] - Local IP address

Selects the connections that match this parameter value.

[-local-port <integer>] - Local Port

Selects the connections that match this parameter value.

[-remote-ip <InetAddress>] - Remote IP Address

Selects the connections that match this parameter value.

[-remote-host <Remote IP>] - Remote Host

Selects the connections that match this parameter value.

[-remote-port <integer>] - Remote Port

Selects the connections that match this parameter value.

[-proto {UDP|TCP|TLS}] - Protocol

Selects the connections that match this parameter value. Possible values are tcp (TCP), udp (UDP), and NA (not applicable).

[-lifid <integer>] - Logical Interface ID

Selects the connections that match this parameter value.

[-service <protocol service>] - Protocol Service

Selects the connections that match this parameter value. Possible values include: nfs, iscsi, and loopback.

[-lru {yes|no}] - Least Recently Used

Selects the connections that match this parameter value.

[-blocks-lb {true|false}] - Connection Blocks Load Balance Migrate

Selects the logical interfaces that are blocked (true) or not blocked (false) from migrating due to an active client connection.

Examples

The following example displays information about active network connections for the node named node0:

```
cluster1::> network connections active show node -node0

Vserver Interface        Remote
Name    Name:Local Port  IP Address:Port   Protocol/Service
------- ---------------- ----------------- ----------------
node0   cluslif1:7070    192.0.2.253:48621 UDP/rclopcp
node0   cluslif1:7070    192.0.2.253:48622 UDP/rclopcp
node0   cluslif2:7070    192.0.2.252:48644 UDP/rclopcp
node0   cluslif2:7070    192.0.2.250:48646 UDP/rclopcp
node0   cluslif1:7070    192.0.2.245:48621 UDP/rclopcp
node0   cluslif1:7070    192.0.2.245:48622 UDP/rclopcp
node0   cluslif2:7070    192.0.2.251:48644 UDP/rclopcp
node0   cluslif2:7070    192.0.2.251:48646 UDP/rclopcp
node0   cluslif1:7070    192.0.2.248:48621 UDP/rclopcp
node0   cluslif1:7070    192.0.2.246:48622 UDP/rclopcp
node0   cluslif2:7070    192.0.2.252:48644 UDP/rclopcp
node0   cluslif2:7070    192.0.2.250:48646 UDP/rclopcp
node0   cluslif1:7070    192.0.2.254:48621 UDP/rclopcp
node0   cluslif1:7070    192.0.2.253:48622 UDP/rclopcp
[...]
```

At privilege levels above "admin", the command displays an extra column.

```
cluster1::*> network connections active show node -node0
                                                                Blocks
Vserver Interface        Remote                                 LB
Name    Name:Local Port  IP Address:Port   Protocol/Service Migrate
------- ---------------- ----------------- ---------------- -------
node0   cluslif1:7070    192.0.2.253:48621 UDP/rclopcp      false
node0   cluslif1:7070    192.0.2.253:48622 UDP/rclopcp      false
node0   cluslif2:7070    192.0.2.252:48644 UDP/rclopcp      false
node0   cluslif2:7070    192.0.2.250:48646 UDP/rclopcp      false
node0   cluslif1:7070    192.0.2.245:48621 UDP/rclopcp      false
node0   cluslif1:7070    192.0.2.245:48622 UDP/rclopcp      false
node0   cluslif2:7070    192.0.2.251:48644 UDP/rclopcp      false
node0   cluslif2:7070    192.0.2.251:48646 UDP/rclopcp      false
node0   cluslif1:7070    192.0.2.248:48621 UDP/rclopcp      false
node0   cluslif1:7070    192.0.2.246:48622 UDP/rclopcp      false
node0   cluslif2:7070    192.0.2.252:48644 UDP/rclopcp      false
node0   cluslif2:7070    192.0.2.250:48646 UDP/rclopcp      false
node0   cluslif1:7070    192.0.2.254:48621 UDP/rclopcp      false
node0   cluslif1:7070    192.0.2.253:48622 UDP/rclopcp      false
[...]
```

network connections listening show

Show the listening connections in this cluster

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

Description

The `network connections listening show` command displays information about network connections that are in an open and listening state.

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>|local}] - Node

Selects the listening connections that match this parameter value.

[-mgmt-cid <integer>] - Management Connection ID

Selects the listening connections that match this parameter value.

[-vserver <vserver>] - Vserver

Selects the listening connections that match this parameter value.

[-cid <integer>] - System Connection ID

Selects the listening connections that match this parameter value.

[-lif-name <lif-name>] - Logical Interface Name

Selects the listening connections that match this parameter value.

[-local-address <IP Address>] - Local IP Address

Selects the listening connections that match this parameter value.

[-local-port <integer>] - Local Port

Selects the listening connections that match this parameter value.

[-remote-ip <InetAddress>] - Remote IP Address

Selects the listening connections that match this parameter value.

[-remote-host <Remote IP>] - Remote Host

Selects the listening connections that match this parameter value.

[-remote-port <integer>] - Remote Port

Selects the listening connections that match this parameter value.

[-proto {UDP|TCP|TLS}] - Protocol

Selects the listening connections that match this parameter value. Possible values include tcp (TCP), udp (UDP), and NA (not applicable).

[-lifid <integer>] - Logical Interface ID

Selects the listening connections that match this parameter value.

[-service <protocol service>] - Protocol Service

Selects the listening connections that match this parameter value. Possible values include: nfs, iscsi, and loopback.

[-lru {yes|no}] - Least Recently Used

Selects the listening connections that match this parameter value.

Examples

The following example displays information about all listening network connections:

```
cluster1::> network connections listening show
Vserver Name Interface Name:Local Port  Protocol/Service
------------ -------------------------- -----------------
node0        cluslif1:7700              UDP/rclopcp
node0        cluslif2:7700              UDP/rclopcp
node1        cluslif1:7700              UDP/rclopcp
node1        cluslif2:7700              UDP/rclopcp
node2        cluslif1:7700              UDP/rclopcp
node2        cluslif2:7700              UDP/rclopcp
node3        cluslif1:7700              UDP/rclopcp
node3        cluslif2:7700              UDP/rclopcp
8 entries were displayed.
```

The following example displays detailed information about listening network connections for the node named node0:

```
cluster1::> network connections listening show -node node0
          Node: node0
Management Connection Id: 0
  System Connection Id: 0
         Vserver: vs0
 Logical Interface Name: datalif1
    Local IP address: 192.0.2.130
       Local Port: 111
    Remote IP address:
       Remote Port: 0
        Protocol: UDP
  Logical Interface Id: 1029
    Protocol Service: port_map
   least recently used: yes
          Node: node0
Management Connection Id: 1
  System Connection Id: 0
         Server: vs0
 Logical Interface Name: datalif2
    Local IP address: 192.0.2.131
       Local Port: 111
    Remote IP address:
       Remote Port: 0
        Protocol: UDP
  Logical Interface Id: 1030
    Protocol Service: port_map
   least recently used: yes
```

network device-discovery commands

network device-discovery show

Display device discovery information

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

Description

The network device-discovery show command displays information about discovered devices. This information may be useful in determining the network topology or investigating connectivity issues. By default, the command displays the following information:

- Local interface
- Discovered device
- Discovered interface
- Discovered platform

Parameters

{ [-fields <fieldname>,…]

Include the specified field or fields in the command output. Use '-fields ?' to display the valid fields.

| [-instance] }

Use this parameter to display detailed information about all fields.

[-node <nodename>] - Node

Displays the discovery ports that match the node name.

[-protocol {cdp|lldp}] - Protocol

Displays the devices that are discovered by the given protocol.

[-port <text>] - Port

Displays the discovery ports that match the physical network port. For example, e0a will display devices discovered on port e0a.

[-discovered-device <text>] - Discovered Device

Displays the discovered devices that match the discovered device name.

[-interface <text>] - Discovered Device Interface

Displays the discovered devices that match this interface port name. The format is dependent on the reporting device. For example: FastEthernet0/12

[-device-ip <IP Address>,…] - Discovered Device IP Addresses

Displays the discovered devices that match the IP address(es). At present, only IPv4 addresses are included. It is recommended to use wildcards around the desired value.

[-platform <text>] - Discovered Device Platform

Displays the discovery ports that contain the platform of discovered devices. For example: N5K-C5010P-BF

[-version <text>] - Discovered Device Version

Displays the discovery ports that contain the version of discovered devices.

[-chassis-id <text>] - Discovered Device Chassis ID

Displays the discovered devices that match the chassis ID.

[-system-name <text>] - Discovered Device System Name

Displays the discovered devices that match the system name.

[-hold-time-remaining <integer>] - Discovered Device's Remaining Hold Time

Displays the discovered devices that match the remaining packet hold time in seconds. If an advertisement from the device isn't received before this time reaches zero, the entry will expire and be removed from the list. For example, "<120" will display discovered devices which will expire within the next 120 seconds.

[-capabilities {router|trans-bridge|source-route-bridge|switch|host|igmp|repeater|phone}] - Discovered Device Capabilities

Displays the discovered devices that match the capability or capabilities. Possible values:

- "router" - Router
- "trans-bridge" - Trans Bridge
- "source-route-bridge" - Source Route Bridge
- "switch" - Switch
- "host" - Host
- "igmp" - IGMP
- "repeater" - Repeater
- "phone" - Phone

Examples

```
cluster1::> network device-discovery show

Node/       Local  Discovered
Protocol    Port   Device                    Interface         Platform
----------- ------ ------------------------- ----------------
----------------
node1/cdp
            e0a    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/17  N5K-C5010P-
BF
            e0b    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/18  N5K-C5010P-
BF
            e1a    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/2   N5K-C5010P-
BF
node2/cdp
            e0a    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/19  N5K-C5010P-
BF
            e0b    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/20  N5K-C5010P-
BF
            e1a    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/4   N5K-C5010P-
BF
            e1c    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/36  N5K-C5010P-
BF
            e1d    US-LS01-5010-F11-NX.example.com(SSI142311PD)
                                             Ethernet100/1/35  N5K-C5010P-
BF
8 entries were displayed.
```

network fcp commands

network fcp adapter modify

Modify the fcp adapter settings

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

Description

Modifies the FCP target adapter information.

The adapter argument is in the form Xy or Xy_z where X and z are integers and y is a letter. An example is 4a or 4a_1.

You cannot bring an adapter offline until all logical interfaces connected to that adapter are offline. Use the network interface modify command to take your logical interfaces offline.

The speed option sets the Fibre Channel link speed of an adapter. You can set adapters that support:

- 10Gb/s to 10 or auto
- 8Gb/s to 2, 4, 8 or auto
- 4Gb/s to 2, 4 or auto
- 2Gb/s to 2 or auto

By default, the link speed option is set to auto for auto negotiation. Setting the link speed to a specific value disables auto negotiation. Under certain conditions, a speed mismatch can prevent the adapter from coming online.

> The system reports the actual link speed with the "Data Link Rate (Gbit)" field in the output of network fcp adapter show`-instance`.

Parameters

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

Specifies the node of the target adapter.

`-adapter <text>` - Adapter

Specifies the target adapter.

`[-status-admin {down|up}]` - Administrative Status

Specifies the desired (administrative) status of the adapter. To view the actual operational status, run network fcp adapter show`-fields`*`status-oper`*.

`[-speed {1|2|4|8|10|16|32|64|auto}]` - Configured Speed

Specifies the adapter configuration speed in Gigabytes.

Examples

```
cluster1::> network fcp adapter modify -node node1 -adapter 0d -speed 2
```

Configures the speed of FCP adapter 0d on node1 to 2 Gb/s.

Related Links

- network interface modify
- network fcp adapter show

network fcp adapter show

Display FCP adapters

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

Description

Displays FCP target adapter information. You can also use this information to determine if adapters are active and online.

The adapter argument is in the form Xy or Xy_z where X and z are integers and y is a letter. An example is 4a or 4a_1.

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>|local}] - Node

If this parameter is specified, the command displays information only about the FCP target adapters that are present on the specified node.

[-adapter <text>] - Adapter

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified name.

[-description <text>] - Description

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified description.

[-physical-protocol {fibre-channel|ethernet}] - Physical Protocol

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified physical protocol. Possible values are *fibre-channel* and *ethernet* .

[-max-speed {1|2|4|8|10|16|32|64|auto}] - Maximum Speed

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified maximum speed.

[-status-admin {down|up}] - Administrative Status

If this parameter is specified, the command displays information only about the FCP target adapters that match the administrative state. Possible values are *up* and *down* .

[-status-oper <text>] - Operational Status

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified operational status.

[-status-extended <text>] - Extended Status

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified extended status.

[-portaddr <Hex Integer>] - Host Port Address

If this parameter is specified, the command displays information only about the FCP target adapters connected with the specified fabric port address.

[-firmware-rev <text>] - Firmware Revision

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified firmware revision.

[-data-link-rate <integer>] - Data Link Rate (Gbit)

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified data link rate.

[-fabric-established {true|false}] - Fabric Established

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified fabric login establishment state.

[-fabric-name <text>] - Fabric Name

If this parameter is specified, the command displays information only about the FCP target adapters that are logged in to the fabric with the specified WWN.

[-conn-established {loop|ptp}] - Connection Established

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified connection type. Possible values are *loop* and *ptp* .

[-is-conn-established {true|false}] - Is Connection Established

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified connection established state.

[-media-type {loop|ptp|auto}] - Mediatype

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified configured media type. Possible values are *loop* , *ptp* , and *auto* .

[-speed {1|2|4|8|10|16|32|64|auto}] - Configured Speed

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified configured speed. If the adapter is set to auto-negotiate, then the value will be *auto* .

[-data-protocols-supported {fcp|fc-nvme}] - Data Protocols Supported

If this parameter is specified, the command displays information only about the FCP target adapters that may host LIFs with the specified data protocol. Possible values are *fcp* and *fc-nvme* .

[-domain-id <integer>] - Domain ID

If this parameter is specified, the command displays information only about the FCP target adapters with a domain identifier that matches the specified domain identifier.

[-fc-wwnn <text>] - Adapter WWNN

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified world wide node name.

[-fc-wwpn <text>] - Adapter WWPN

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified world wide port name.

[-switch-port <text>] - Switch Port

If this parameter is specified, the command displays information only about the FCP target adapters that are connected to the specified switch port.

[-sfp-formfactor <text>] - Form Factor Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP form factor.

[-sfp-vendor-name <text>] - Vendor Name Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP vendor name.

[-sfp-part-number <text>] - Part Number Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP part number.

[-sfp-rev <text>] - Revision Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP revision number.

[-sfp-serial-number <text>] - Serial Number Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP serial number.

[-sfp-fc-speed-capabilities <text>] - FC Capabilities Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP speed capabilities.

[-sfp-vendor-oui <text>] - Vendor OUI Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP vendor OUI.

[-sfp-wavelength <integer>] - Wavelength In Nanometers

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP wavelength.

[-sfp-date-code <text>] - Date Code Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP date code.

[-is-sfp-optical-transceiver-valid {true|false}] - Validity Of Transceiver

If this parameter is specified, the command displays information only about the FCP target adapters that

match whether the SFP is installed and valid.

[-sfp-connector <text>] - Connector Used

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP connector type.

[-sfp-encoding <text>] - Encoding Used

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP physical encoding.

[-is-sfp-diagnostics-internally-calibrated {true|false}] - Is Internally Calibrated

If this parameter is specified, the command displays information only about the FCP target adapters that match whether the SFP diagnostics are internally calibrated or not.

[-sfp-diagnostic-monitoring-type <Hex Integer>] - Diagnostic Monitoring Type

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP diagnostic monitoring type.

[-sfp-ddm-capabilities <text>] - Status Monitoring Available

If this parameter is specified, the command displays information only about the FCP target adapters that match whether the specified SFP digital diagnostics monitoring are supported or not.

[-sfp-sff8472-compliance <Hex Integer>] - SFF-8472 Compliance

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP SFF8472 compliance.

[-sfp-rx-power <text>] - Received Optical Power

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified observed SFP receive power.

[-is-sfp-rx-power-in-range {true|false}] - Is Received Power In Range

If this parameter is specified, the command displays information only about the FCP target adapters that match whether the observed SFP receive power is within the valid range for the SFP.

[-sfp-tx-power <text>] - SFP Transmitted Optical Power

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP transmit power.

[-is-sfp-tx-power-in-range {true|false}] - Is Xmit Power In Range

If this parameter is specified, the command displays information only about the FCP target adapters that match whether the observed SFP transmit power is within the valid range for the SFP.

[-sfp-ddm-status-control <Hex Integer>] - DDM Status

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified SFP DDM status and control.

[-is-sfp-tx-in-disable {true|false}] - Is Xmit Disabled

If this parameter is specified, the command displays information only about the FCP target adapters that match whether the observed SFP transmitter is in disabled state.

[-is-sfp-tx-in-fault {true|false}] - Is Xmit In Fault

If this parameter is specified, the command displays information only about the FCP target adapters that match whether the observed SFP transmitter is in faulted state.

[-is-sfp-rx-in-los {true|false}] - Is Receiver In LOS

If this parameter is specified, the command displays information only about the FCP target adapters that match whether the observed SFP receiver is in loss of signal state.

Examples

```
cluster1::> fcp adapter show
                     Connection  Host
Node         Adapter Established Port Address
------------ ------- ----------- ------------
sti6280-021   0a      ptp         30012c
```

The example above displays information regarding FCP adapters within cluster1.

```
cluster1::> fcp adapter show -instance -node sti6280-021 -adapter 0a
Node: sti6280-021
                      Adapter: 0a
                  Description: Fibre Channel Target Adapter 0a (QLogic
2532 (2562), rev. 2, 8G)
             Physical Protocol: fibre-channel
                 Maximum Speed: 8
         Administrative Status: up
            Operational Status: online
               Extended Status: ADAPTER UP
             Host Port Address: 30012c
             Firmware Revision: 5.8.0
         Data Link Rate (Gbit): 4
            Fabric Established: true
                   Fabric Name: 20:14:54:7f:ee:54:b9:01
        Connection Established: ptp
     Is Connection Established: true
                     Mediatype: ptp
              Configured Speed: auto
                  Adapter WWNN: 50:0a:09:80:8f:7f:8b:1c
                  Adapter WWPN: 50:0a:09:81:8f:7f:8b:1c
                   Switch Port: RTP-AG01-410B51:1/41
    Form Factor Of Transceiver: SFP
    Vendor Name Of Transceiver: OPNEXT,INC
    Part Number Of Transceiver: TRS2000EN-SC01
       Revision Of Transceiver: 0000
  Serial Number Of Transceiver: T10H64793
```

```
 FC Capabilities Of Transceiver: 10 (Gbit/sec)
      Vendor OUI Of Transceiver: 0:11:64
       Wavelength In Nanometers: 850
       Date Code Of Transceiver: 10:08:17
        Validity Of Transceiver: true
                 Connector Used: LC
                  Encoding Used: 64B66B
       Is Internally Calibrated: true
     Diagnostic Monitoring Type: 68
    Status Monitoring Available: fa {Rx_Loss_of_Sig, Tx_Fault, Tx_Disable}
            SFF-8472 Compliance: 5
         Received Optical Power: 441.3 (uWatts)
     Is Received Power In Range: true
  SFP Transmitted Optical Power: 600.4 (uWatts)
        Is Xmit Power In Range: true
                    DDM Status: 30
              Is Xmit Disabled: false
              Is Xmit In Fault: false
            Is Receiver In LOS: false
```

The example above displays detailed information regarding FCP adapter 0a in sti6280-021 within cluster1.

network fcp topology show

FCP topology interconnect elements per adapter

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

Description

Display FCP topology interconnect elements per adapter.

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>|local}] - Node

Use this parameter to select the interconnect elements for adapters that are located on the node that you specify.

[-adapter <text>] - Adapter

Use this parameter to select the interconnect elements for the specified adapter.

[-domain-id <integer>] - Domain Id

Use this parameter to select the interconnect elements with the specified domain identifier.

[-port-wwpn <text>] - Port WWPN

Use this parameter to select the interconnect elements with the specified port world wide name.

[-switch-name <text>] - Switch Name

Use this parameter to select the interconnect elements with the specified switch.

[-switch-vendor <text>] - Switch Vendor

Use this parameter to select the interconnect elements with the specified vendor.

[-switch-release <text>] - Switch Release

Use this parameter to select the interconnect elements with the specified release.

[-switch-wwpn <text>] - Switch WWPN

Use this parameter to select the interconnect elements with the specified world wide port name.

[-switch-wwn <text>] - Switch WWN

Use this parameter to select the interconnect elements with the specified world wide name.

[-port-count <integer>] - Port Count

Use this parameter to select the interconnect elements with the specified port count.

[-port-slot <text>] - Port Slot

Use this parameter to select the interconnect elements with the specified port slot.

[-port-state {Unknown|Online|Offline|Testing|Fault}] - Port State

Use this parameter to select the interconnect elements with the specified port state.

[-port-type {None|N-Port|NL-Port|FNL-Port|NX-Port|F-Port|FL-Port|E-Port|B-Port|TNP-Port|TF-Port|NV-Port|FV-Port|SD-Port|TE-Port|TL-Port}] - Port Type

Use this parameter to select the interconnect elements with the specified port type.

[-port-attached-wwpn <text>] - Attached Port WWPN

Use this parameter to select the interconnect elements with the specified attached wwpn.

[-port-attached-id <text>] - Attached Port Id

Use this parameter to select the interconnect elements with the specified attached id.

[-port-attached-visible <text>] - Visible

Use this parameter to select the interconnect elements with the specified visibility flag on attached port structure.

Examples

```
cluster1::> network fcp topology show
Switch connected to the adapter 0c
      Switch Name: ssan-fc0e-d58
    Switch Vendor: Cisco Systems, Inc.
   Switch Release: 5.2(1)N1(9)
    Switch Domain: 4
       Switch WWN: 20:05:00:05:9b:26:f4:c1
       Port Count: 20

Port    Port WWN                State     Type      Attached WWPN
Port Id
------- ----------------------- --------- ---------
----------------------- ------------
vfc9    20:08:00:05:9b:26:f4:ff Offline   None      -
-
vfc10   20:15:00:05:9b:26:f4:ff Online    TF-Port
50:0a:09:82:8d:92:4c:ff 0x0407c0  *
vfc11   20:16:00:05:9b:26:f4:ff Online    TF-Port
50:0a:09:81:8d:e2:4e:ec 0x040800  *

Switch connected to the adapter 0c
      Switch Name: ssan-fc0e-d58
    Switch Vendor: Cisco Systems, Inc.
   Switch Release: 5.2(1)N1(9)
    Switch Domain: 4
       Switch WWN: 20:05:00:05:9b:26:f4:c1
       Port Count: 20

Port    Port WWN                State     Type      Attached WWPN
Port Id
------- ----------------------- --------- ---------
----------------------- ------------
vfc20   20:13:00:05:9b:26:f4:ff Offline   None      -
-
vfc21   20:14:00:05:9b:26:f4:ff Online    TF-Port
50:0a:09:81:8d:92:4c:ff 0x0407a0  *

5 entries were displayed.
```

The example above show FCP topology interconnect information for the cluster.

network fcp zone show

Display the active zone set information

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

Description

Displays the active zone set information.

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>|local}] - Node

Use this parameter to select the active zone set information for adapters that located on the node name that you specify.

[-adapter <text>] - Adapter

Use this parameter to select the active zone set information for the specified adapter.

[-zoneset-name <text>] - Zoneset Name

Use this parameter to select the active zone set information for the specified zone set name.

[-zone-name <text>] - Zone Name

Use this parameter to select the active zone set information for the specified zone name.

[-type-name <text>] - Type Name

Use this parameter to select the active zone set information with the specified symbolic type.

[-type <integer>] - Type

Use this parameter to select the active zone set information with the specified port type.

[-port-id <Hex Integer>] - Member Port Id

Use this parameter to select the active zone set information with the specified member port id.

[-domain-id <integer>] - Member Domain Id

Use this parameter to select the active zone set information with the specified member domain id.

[-port <integer>] - Member Port

Use this parameter to select the active zone set information with the specified member port.

[-wwn <text>] - Member WWN

Use this parameter to select the active zone set information with the specified member WWN.

[-zone-count <integer>] - Zone Count

Use this parameter to select the active zone set information with the specified number of zones.

[-zone-member-count <integer>] - Zone Member Count

Use this parameter to select the active zone set information with the specified number of zone members in a zone.

[-contents <text>] - Member Contents

Use this parameter to select the active zone set information using any type.

Examples

```
cluster1::> network fcp adapter zone show
                                            Member
                       Zone Name            Type                WWN
 --------------------- ------------------ -------------------
-----------------
Active Zone Set on adapter 0c
   Zone Set Name: zoneset_name
                       zone_name_1         Port ID             -
                       zone_name_1         Port ID             -
                       zone_name_1         Port ID             -
                       zone_name_2         Domain ID/Port      -
                       zone_name_2         Domain ID/Port      -
                       zone_name_2         Domain ID/Port      -
                       zone_name_3         Fabric Port Name
00:00:00:00:00:00:00:00
                       zone_name_3         Fabric Port Name
01:00:00:00:00:00:00:00
                       zone_name_3         Fabric Port Name
02:00:00:00:00:00:00:00

9 entries were displayed.
```

The example above displays information regarding active zone set information for the cluster.

network interface commands

network interface create

Create a logical interface

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

Description

The `network interface create` command creates a logical interface (LIF).

> A logical interface is an IP address associated with a physical network port. For logical interfaces using NAS data protocols, the interface can fail over or be migrated to a different physical port in the event of component failures, thereby continuing to provide network access despite the component failure.

On some cloud platforms, this operation might perform changes to the external route tables.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the Vserver on which the LIF is created.

`-lif <lif-name>` - Logical Interface Name

Use this parameter to specify the name of the LIF that is created. For iSCSI and FC LIFs, the name cannot be more than 254 characters.

`[-service-policy <text>]` - Service Policy

Use this parameter to specify a service policy for the LIF. If no policy is specified, a default policy will be assigned automatically. Use the network interface service-policy show command to review available service policies.

`[-role {undef|cluster|data|node-mgmt|intercluster|cluster-mgmt}]` - (DEPRECATED)-Role

This parameter has been deprecated and may be removed in a future version of ONTAP. Use the `-service-policy` parameter instead.

Use this parameter to specify the role of the LIF. LIFs can have one of five roles:

- Cluster LIFs, which provide communication among the nodes in a cluster
- Intercluster LIFs, which provide communication among peered clusters
- Data LIFs, which provide data access to NAS and SAN clients
- Node-management LIFs, which provide access to cluster management functionality
- Cluster-management LIFs, which provide access to cluster management functionality

LIFs with the cluster-management role behave as LIFs with the node-management role except that cluster-management LIFs can failover between nodes.

`[-data-protocol {nfs|cifs|iscsi|fcp|fcache|none|fc-nvme|s3|nvme-tcp}]` - Data Protocol

Use this parameter to specify the list of data protocols that can be served by the LIF. The supported protocols are NFS, CIFS, iSCSI, FCP, and FC-NVMe. NFS and CIFS are available by default when you create a LIF. If you specify "none", the LIF does not support any data protocols. Also, none, iscsi, fcp or fc-nvme cannot be combined with any other protocols.

The data-protocol field must be specified when the LIF is created and cannot be modified later.

The NFS protocol relies on firewall services included in the built-in "data" and "mgmt-nfs" firewall policies. Assigning a different firewall policy might disrupt some NFS client implementations.

`-address <IP Address>` - Network Address

Use this parameter to specify the LIF's IP address.

A cluster LIF cannot be on the same subnet as a management or data LIF.

{ -netmask <IP Address> - Netmask

Use this parameter to specify the LIF's netmask.

| -netmask-length <integer> - Bits in the Netmask

Use this parameter to specify the length (in bits) of the LIF's netmask.

| -is-vip <true> - Is VIP LIF

Use this parameter to display only logical interfaces matching a specify "is-vip" flag. Specifying "true" matches only LIFs to implement a Virtual IP; "false" matches only LIFs that do not.

{ [-auto <true>] - Allocate Link Local IPv4 Address

Use this parameter to specify whether IPv4 link local addressing is enabled for this LIF.

| [-subnet-name <subnet name>] - Subnet Name }

Use this parameter to allocate the interface address from a subnet. If needed, a default route will be created for this subnet.

[-home-node <nodename>] - Home Node

Use this parameter to specify the LIF's home node. The home node is the node to which the LIF returns when the network interface revert command is run on the LIF.

[-home-port {<netport>|<ifgrp>}] - Home Port

Use this parameter to specify the LIF's home port or interface group. The home port is the port or interface group to which the LIF returns when the network interface revert command is run on the LIF.

[-status-admin {up|down}] - Administrative Status

Use this parameter to specify whether the initial administrative status of the LIF is up or down. The default setting is `up` . The administrative status can differ from the operational status For example, if you specify the status as up but a network problem prevents the interface from functioning, the operational status remains as down.

[-failover-policy {system-defined|local-only|sfo-partner-only|disabled|broadcast-domain-wide}] - Failover Policy

Use this parameter to specify the failover policy for the LIF.

- system-defined - The system determines appropriate failover targets for the LIF. The default behavior is that failover targets are chosen from the LIF's current hosting node and also from one other non-partner node when possible.
- local-only - The LIF fails over to a port on the local or home node of the LIF.
- sfo-partner-only - The LIF fails over to a port on the home node or SFO partner only.
- broadcast-domain-wide - The LIF fails over to a port in the same broadcast domain as the home port.
- disabled - Failover is disabled for the LIF.

The failover policy for cluster logical interfaces is local-only and cannot be changed. The default failover policy for data logical interfaces is system-defined. This value can be changed.

[-firewall-policy <policy>] - (DEPRECATED)-Firewall Policy

> This parameter has been deprecated and may be removed in a future version of ONTAP. Use the `-service-policy` parameter instead.

Use this parameter to specify the firewall policy for the LIF. A LIF can use a default firewall policy that corresponds to its role (management, cluster, intercluster, or data) or a custom firewall policy created by an administrator. View and modify existing firewall policies using the system services firewall policy show and system services firewall policy modify commands, respectively.

> The NFS data protocol relies on firewall services included in the built-in "data" and "mgmt-nfs" firewall policies. Assigning a different firewall policy might disrupt some NFS client implementations.

[-auto-revert {true|false}] - Auto Revert

Use this parameter to specify whether a data LIF is automatically reverted to its home port under certain circumstances. These circumstances include startup, when the status of the management database changes to either master or secondary, or when the network connection is made.

[-dns-zone {<zone-name>|none}] - Fully Qualified DNS Zone Name

Use this parameter to specify a unique, fully qualified domain name of a DNS zone to which this data LIF is added. You can associate a data LIF with a single DNS zone. All data LIFs included in a zone must be on the same Vserver. If a LIF is not added to a DNS zone the data LIF is created with the value `none` .

[-listen-for-dns-query {true|false}] - DNS Query Listen Enable

Use this parameter to specify if the LIF has to listen for DNS queries. The default value for this parameter is true.

[-allow-lb-migrate {true|false}] - (DEPRECATED)-Load Balancing Migrate Allowed (privilege: advanced)

> This parameter has been deprecated and may be removed in a future version of ONTAP.

Use this parameter to specify whether load balancing migration is activated for this data LIF. The default value of this parameter is `false` . If you set the value of this parameter to `true` , automatic revert capability for this data LIF is disabled (the `-auto-revert` parameter is set to `false`). Also, data LIFs that migrate as a result of load balancing adhere to network interface failover rules.

> During times when a LIF is hosting active NFSv4, CIFS, or NRV connections, load balancing based LIF migrations between nodes will be temporarily disabled.

[-lb-weight {load|0..100}] - Load Balanced Weight (privilege: advanced)

Use this parameter to specify a load balancing weight for a data LIF. A valid numeric load balancing weight is any integer between 0 and 100. When you specify the same load balancing weight for all data LIFs in a DNS zone, client requests are uniformly distributed, similar to round-robin DNS. A data LIF with a low load balancing weight is made available for client requests less frequently than one that has a high load balancing weight. "load" is the default value of this parameter. If set to "load", node utilization statistics are used to dynamically assign the load balancing weight.

[-failover-group <failover-group>] - Failover Group Name

Use this parameter to specify the name of the failover group to associate with the LIF. Manage failover groups by using the `network interface failover-groups` command. Each broadcast domain has a

default failover group which is created by the system automatically and has the same name as the broadcast domain. The failover group associated with the broadcast domain includes all ports in the broadcast domain. A logical interface's failover group is set to the failover group of the home port's broadcast domain by default, but this value can be modified.

[-comment <text>] - Comment

Use this parameter to specify the comment to associate with the LIF.

[-force-subnet-association <true>] - Force the LIF's Subnet Association

This command will fail if the IP address falls within the address range of a named subnet. Set this to true to acquire the address from the named subnet and assign the subnet to the LIF.

[-is-dns-update-enabled {true|false}] - Is Dynamic DNS Update Enabled?

If this parameter is set to *true* , then dynamic DNS update is sent to the DNS server for the particular LIF entry if dynamic DNS updates are enabled for the corresponding Vserver. This field is set to *true* by default for both IPv4 and IPv6 LIFs. DNS Update is not supported on LIFs not configured with either the NFS or CIFS protocol.

[-probe-port <integer>] - Probe-port for Cloud Load Balancer

Use this parameter to specify a probe-port for the LIF in the Azure environment. It is a required field in the Azure environment. If no probe-port is specified, an error would be returned.

[-broadcast-domain <text>] - Broadcast Domain

Use this parameter to display the broadcast domain that contains the home port of the logical interface.

[-rdma-protocols <roce>,…] - Required RDMA offload protocols

Defines RDMA offload protocols required by the LIF. A non-empty list will ensure that this LIF can only be moved to network ports that support the specified RDMA offload protocols.

Examples

The following example creates an IPv4 LIF named datalif1 and an IPv6 LIF named datalif2 on a Vserver named vs0. Their home node is node0 and home port is e0c. The failover policy *broadcast-domain-wide* is assigned to both LIFs. The service policy is *default-data-files* and the LIFs are automatically reverted to their home node at startup and under other circumstances. The datalif1 has the IP address 192.0.2.130 and netmask 255.255.255.128, and datalif2 has the IP address 3ffe:1::aaaa and netmask length of 64.

```
cluster1::> network interface create -vserver vs0 -lif datalif1 -home-node
node0 -home-port e0c -address 192.0.2.130 -netmask 255.255.255.128
-failover-policy broadcast-domain-wide -service-policy default-data--files
-auto-revert true
cluster1::> network interface create -vserver vs0 -lif datalif2 -home-node
node0 -home-port e0c -address 3ffe:1::aaaa -netmask-length 64 -failover
-policy broadcast-domain-wide -service-policy default-data-files -auto
-revert true
```

Related Links

- network interface service-policy show

- network interface revert
- system services firewall policy show
- system services firewall policy modify

network interface delete

Delete a logical interface

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

Description

The `network interface delete` command deletes a logical interface (LIF) from a Vserver. Only administratively down LIFs can be deleted. To make a LIF administratively down, use the network interface modify command to set the "status-admin" parameter to "down".

> If the LIF is configured for a SAN protocol and is part of a portset, the LIF must be removed from the portset before it can be deleted. To determine if a LIF is in a portset, use the lun portset show command. To remove the LIF from the portset, use the lun portset remove command.

> On some cloud platforms, this operation might perform changes to the external route tables.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the Vserver on which the logical interface to be deleted is located.

`-lif <lif-name>` - Logical Interface Name

Use this parameter to specify the logical interface to delete.

Examples

The following example deletes a logical interface named cluslif3 that is located on a Vserver named vs0.

```
cluster1::> network interface delete -vserver vs0 -lif cluslif3
```

Related Links

- network interface modify
- lun portset show
- lun portset remove

network interface migrate-all

Migrate all data logical interfaces away from the specified node

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

Description

The `network interface migrate-all` command migrates all data logical interfaces from the node you specify.

> Manual migration of a logical interface can take up to 15 seconds to complete. Logical interface migration is a best-effort command and can only be completed if the destination node and port are operational. Logical interface migration requires that the logical interface be pre-configured with valid failover rules to facilitate failover to a remote node.

> On some cloud platforms, this operation might perform changes to the external route tables.

Parameters

`-node <nodename>` - Node

Use this parameter to specify the node from which all logical interfaces are migrated. Each data logical interface is migrated to another node in the cluster, assuming that the logical interface is configured with failover rules that specify an operational node and port.

`[-port {<netport>|<ifgrp>}]` - Port

Use this parameter to specify the port from which all logical interfaces are migrated. This option cannot be used with asynchronous migrations. If this parameter is not specified, then logical interfaces will be migrated away from all ports on the specified node.

Examples

The following example migrates all data logical interfaces from the current (local) node.

```
cluster1::> network interface migrate-all -node local
```

network interface migrate

Migrate a logical interface to a different port

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

Description

The `network interface migrate` command migrates a logical interface to a port or interface group on the node you specify.

> Manual migration of a logical interface can take up to 15 seconds to complete. Also, when you migrate a cluster logical interface, you must do so from the local node. Logical interface migration is a best-effort command, and can only be completed if the destination node and port are operational

> On some cloud platforms, this operation might perform changes to the external route tables.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the Vserver that owns the logical interface that is to be migrated.

`-lif <lif-name>` - Logical Interface Name

Use this parameter to specify the logical interface that is to be migrated.

`-destination-node <nodename>` - Destination Node

Use this parameter to specify the node to which the logical interface is to be migrated.

`[-destination-port {<netport>|<ifgrp>}]` - Destination Port

Use this parameter to specify the port or interface group to which the logical interface is to be migrated.

`[-force <true>]` - Force Migrate Data LIF Flag (privilege: advanced)

Use this parameter to force the migration operation. Intended for disaster recovery only. Specifying this parameter implies the LIF's current node is down and will skip removing the LIF from its current node.

Examples

The following example migrates a logical interface named datalif1 on a Vserver named vs0 to port e0c on a node named node2:

```
cluster1::> network interface migrate -vserver vs0 -lif datalif1 -dest
-node node2 -dest-port e0c
```

network interface modify

Modify a logical interface

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

Description

The `network interface modify` command modifies attributes of a logical interface (LIF).

> You cannot modify some properties of an iSCSI or FCP LIF, such as `-home-node` or `-home-port` , if the LIF is in a portset. To modify these properties, first remove the LIF from the portset. To determine if a LIF is in a portset, use the lun portset show command. To remove the LIF from the portset, use the lun portset remove command.

> On some cloud platforms, this operation might perform changes to the external route tables.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the Vserver on which the LIF to be modified is located.

`-lif <lif-name>` - Logical Interface Name

Use this parameter to specify the name of the LIF that is to be modified

`[-service-policy <text>]` - Service Policy

Use this parameter to modify the service policy associated with the LIF.

`[-address <IP Address>]` - Network Address

Use this parameter to modify the LIF's IP address.

> A cluster LIF cannot be on the same subnet as a data or management LIF.

{ `[-netmask <IP Address>]` - Netmask

Use this parameter to modify the LIF's netmask.

| `[-netmask-length <integer>]` - Bits in the Netmask

Use this parameter to modify the length (in bits) of the LIF's netmask.

| `[-subnet-name <subnet name>]` - Subnet Name }

Use this parameter to allocate the interface address from a subnet. Modifying this parameter will cause a new IP address to be allocated and assigned to the interface.

`[-home-node <nodename>]` - Home Node

Use this parameter to modify the LIF's home node. The home node is the node to which the LIF returns when the network interface revert command is run on that LIF.

`[-home-port {<netport>|<ifgrp>}]` - Home Port

Use this parameter to modify the LIF's home port. The home port is the port or interface group to which the LIF returns when the network interface revert command is run on that LIF.

`[-status-admin {up|down}]` - Administrative Status

Use this parameter to modify the administrative status of the LIF. The administrative status can differ from the operational status. For example, if you specify the status as `up` but a network problem prevents the interface from functioning, the operational status remains as `down` .

`[-failover-policy {system-defined|local-only|sfo-partner-only|disabled|broadcast-domain-wide}]` - Failover Policy

Use this parameter to modify the failover policy for the LIF.

- system-defined - The system determines appropriate failover targets for the LIF. The default behavior is that failover targets are chosen from the LIF's current hosting node and also from one other non-partner node when possible.
- local-only - The LIF fails over to a port on the local or home node of the LIF.
- sfo-partner-only - The LIF fails over to a port on the home node or SFO partner only.
- broadcast-domain-wide - The LIF fails over to a port in the same broadcast domain as the home port.
- disabled - Failover is disabled for the LIF.

> The failover policy for cluster logical interfaces is local-only and cannot be changed. The default failover policy for data logical interfaces is system-defined. This value can be changed.

`[-firewall-policy <policy>]` - (DEPRECATED)-Firewall Policy

> This parameter has been deprecated and may be removed in a future version of ONTAP. Use the `-service-policy` parameter instead.

Use this parameter to set the firewall policy for the LIF. A LIF can use a default firewall policy that corresponds to its role (management, cluster, or data) or a custom firewall policy created by an administrator. When using a custom policy, the interface will fallback on its role's default policy for unspecified services. View existing firewall policies with the "system services firewall policy show " command. Modify existing firewall policies with the "system services firewall policy modify " command.

> The NFS data protocol relies on firewall services included in the built-in "data" and "mgmt-nfs" firewall policies. Assigning a different firewall policy policy might disrupt some NFS client implementations.

`[-auto-revert {true|false}]` - Auto Revert

Use this parameter to modify whether a data LIF is reverted automatically to its home port under certain circumstances. These circumstances would include startup, when the status of the management database changes to either master or secondary, and when the network connection is made.

`[-dns-zone {<zone-name>|none}]` - Fully Qualified DNS Zone Name

Use this parameter to modify the unique, fully qualified domain name of the DNS zone to which this data LIF belongs. You can associate a data LIF with a single DNS zone. All data LIFs included in a zone must be on the same Vserver. If you do not specify a value for this parameter, the data LIF is created with the value `none` .

`[-listen-for-dns-query {true|false}]` - DNS Query Listen Enable

Use this parameter to specify if the LIF has to listen for DNS queries. The default value for this parameter is true.

`[-allow-lb-migrate {true|false}]` - (DEPRECATED)-Load Balancing Migrate Allowed (privilege: advanced)

> This parameter has been deprecated and may be removed in a future version of ONTAP.

Use this parameter to modify whether or not load balancing migration is enabled for this data LIF. The default value of this parameter is `false` . If you set the value of this parameter to `true` , the automatic revert capability of the data LIF is disabled (the `-auto-revert` parameter is set to `false`). Also, data LIFs that migrate as a result of load balancing adhere to network interface failover rules.

> During times when a LIF is hosting active NFSv4, CIFS, or NRV connections, load balancing based LIF migrations between nodes will be temporarily disabled.

`[-lb-weight {load|0..100}]` - Load Balanced Weight (privilege: advanced)

Use this parameter to modify the load balancing weight of the data LIF. A valid load balancing weight is any integer between 1 and 100. If you specify the same load balancing weight for all data LIFs in a DNS zone, client requests are uniformly distributed, similar to round-robin DNS. A data LIF with a low load balancing

weight is made available for client requests less frequently than one that has a high load balancing weight.

[-failover-group <failover-group>] - Failover Group Name

Use this parameter to modify the name of the failover group to associate with the network interface. Manage failover groups using the `network interface failover-groups` command. Each broadcast domain has a default failover group which is created by the system automatically and has the same name as the broadcast domain. The failover group associated with the broadcast domain includes all ports in the broadcast domain. A logical interface's failover group is set to the failover group of the home port's broadcast domain by default, but this value can be modified.

[-comment <text>] - Comment

Use this parameter to modify the comment associated with the LIF.

[-force-subnet-association <true>] - Force the LIF's Subnet Association

This command will fail if the IP address falls within the address range of a named subnet. Set this to true to acquire the address from the named subnet and assign the subnet to the LIF.

[-is-dns-update-enabled {true|false}] - Is Dynamic DNS Update Enabled?

If this parameter is set to *true* , then dynamic DNS update is sent to the DNS server for the particular LIF entry if dynamic DNS updates are enabled for the corresponding Vserver. This field is set to *true* by default for both IPv4 and IPv6 LIFs. DNS Update is not supported on LIFs not configured with either the NFS or CIFS protocol.

[-rdma-protocols <roce>,…] - Required RDMA offload protocols

Defines RDMA offload protocols required by the LIF. A non-empty list will ensure that this LIF can only be moved to network ports that support the specified RDMA offload protocols.

Examples

The following example modifies a LIF named datalif1 on a logical server named vs0. The LIF's netmask is modified to 255.255.255.128.

```
cluster1::> network interface modify -vserver vs0 -lif datalif1 -netmask
255.255.255.128
```

Related Links

- lun portset show
- lun portset remove
- network interface revert
- system services firewall policy show
- system services firewall policy modify

network interface rename

Rename a logical interface

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

Description

Use the `network interface rename` command to change the name of an existing logical interface.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the Vserver on which the logical interface to rename is located.

`-lif <lif-name>` - Logical Interface Name

Use this parameter to specify the name of the logical interface to rename.

`-newname <text>` - The new name for the interface

Use this parameter to specify the new name of the logical interface. For iSCSI and FC LIFs, the name cannot be more than 254 characters.

Examples

The following example renames a cluster logical interface named cluslif1 to cluslif4 on a Vserver named vs0.

```
cluster1::> network interface rename -vserver vs0 -lif cluslif1 -newname
cluslif4
```

network interface revert

Revert a logical interface to its home port

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

Description

The `network interface revert` command reverts a logical interface that is not currently on its home port to its home port, assuming that the home node and port are both operational. A logical interface's home port is specified when the logical interface is created. Determine a logical interface's home port by using the network interface show command.

When you revert a cluster logical interface, you must do so from the local node.

On some cloud platforms, this operation might perform changes to the external route tables.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the Vserver on which the logical interface to be reverted is located.

`-lif <lif-name>` - Logical Interface Name

Use this parameter to specify the logical interface that is to be reverted.

> Logical interfaces for SAN protocols are always home. Thus, this command has no effect on such interfaces. The same applies to logical interfaces for NAS protocols that are already home.

Examples

The following example returns any logical interfaces that are not currently on their home ports to their home ports.

```
cluster1::> network interface revert -vserver * -lif *
```

Related Links

- network interface show

network interface show

Display logical interfaces

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

Description

The `network interface show` command displays information about logical interfaces.

Running the command with the `-failover` parameter displays information relevant to logical interface failover rules.

Running the command with the `-status` parameter displays information relevant to logical interface operational status.

Running the command with the `-by-ipspace` parameter displays information relevant to logical interfaces on a specific IPspace.

See the examples for more information.

You can specify additional parameters to display only information that matches those parameters. For example, to display information only about logical interfaces whose operational status is down, run the command with the `-status-oper down` parameter.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| [-by-ipspace]

Use this parameter to display logical-interfaces sorted by IPspace and Vserver.

| [-dns-zones]

Use this parameter to display logical-interfaces and whether the interface is associated with a Domain Name System (DNS) load balancing zone.

| [-failover]

Use this parameter to display logical-interface failover information.

| [-status]

Use this parameter to display detailed logical-interface status information.

| [-instance] }

Use this parameter to display all the fields for the specified logical-interfaces.

[-vserver <vserver>] - Vserver Name

Use this parameter to display information only about logical interfaces on the Vserver you specify.

Use this parameter plus the `-lif` parameter to display detailed information only about the logical interface you specify.

[-lif <lif-name>] - Logical Interface Name

Use this parameter to display information only about logical interfaces that match the name you specify.

Use this parameter with the `-vserver` parameter to display detailed information only about the logical interface you specify.

[-service-policy <text>] - Service Policy

Use this parameter to display information only about logical interfaces that have the service policy you specify.

[-services <LIF Service Name>,…] - Service List

Use this parameter to display information only about logical interfaces that support all services in a comma-separated list of service names.

[-role {undef|cluster|data|node-mgmt|intercluster|cluster-mgmt}] - (DEPRECATED)-Role

> This parameter has been deprecated and may be removed in a future version of ONTAP. Use either the `-service-policy` or `-services` parameter instead.

Use this parameter to display information only about logical interfaces that are associated with network ports that have the role you specify.

[-data-protocol {nfs|cifs|iscsi|fcp|fcache|none|fc-nvme|s3|nvme-tcp}] - Data Protocol

Use this parameter to display information only about logical interfaces that have the enabled data protocols you specify.

[-address <IP Address>] - Network Address

Use this parameter to display information only about logical interfaces that match the IP address or address range you specify.

[-netmask <IP Address>] - Netmask

Use this parameter to display information only about logical interfaces that have the netmask you specify.

[-netmask-length <integer>] - Bits in the Netmask

Use this parameter to display information only about logical interfaces with a netmask that has the number of bits you specify.

[-is-vip <true>] - Is VIP LIF

Use this parameter to display information only about logical interfaces that are VIP LIFs or not as you specify.

[-subnet-name <subnet name>] - Subnet Name

Use this parameter to display the logical interfaces that matches the subnet name.

[-home-node <nodename>] - Home Node

Use this parameter to display information only about logical interfaces that have the home node you specify.

[-home-port {<netport>|<ifgrp>}] - Home Port

Use this parameter to display information only about logical interfaces that have the home port or interface group you specify.

[-curr-node <nodename>] - Current Node

Use this parameter to display information only about logical interfaces that are currently located on the node you specify.

[-curr-port {<netport>|<ifgrp>}] - Current Port

Use this parameter to display information only about logical interfaces that are currently located on the port or interface group you specify.

[-status-oper {up|down}] - Operational Status

Use this parameter to display information only about logical interfaces that have the operational status you specify.

[-status-extended <text>] - Extended Status

Use this parameter to display information only about logical interfaces that match the extended status that you specify.

[-numeric-id <integer>] - Numeric ID (privilege: advanced)

Use this parameter to display information only about logical interfaces with the numeric ID (or range of IDs) you specify. The numeric ID is an integer that identifies the logical interface in the cluster.

[-is-home {true|false}] - Is Home

Use this parameter to display information only about logical interfaces that are (true) or are not (false) currently located on their home node and port.

[-status-admin {up|down}] - Administrative Status

Use this parameter to display information only about logical interfaces that have the administrative status you specify.

[-failover-policy {system-defined|local-only|sfo-partner-only|disabled|broadcast-domain-wide}] - Failover Policy

Use this parameter to display information only about logical interfaces that use the failover policy you specify.

[-firewall-policy <policy>] - (DEPRECATED)-Firewall Policy

> This parameter has been deprecated and may be removed in a future version of ONTAP. Use the `-service-policy` parameter instead.

Use this parameter to display information only about logical interfaces that use the firewall policies you specify.

[-auto-revert {true|false}] - Auto Revert

Use this parameter to display information only about logical interfaces that have auto-revert setting you specify.

[-sticky {true|false}] - Sticky Flag (privilege: advanced)

Use this parameter to display information only about logical interfaces that are "sticky". A sticky logical interface is one that has been manually migrated to another node and is not subject to auto-revert settings. A sticky logical interface remains at the migrated location until it is manually reverted or until it fails over to another node.

[-dns-zone {<zone-name>|none}] - Fully Qualified DNS Zone Name

Use this parameter to display information only about logical interfaces in the specified DNS zone.

[-listen-for-dns-query {true|false}] - DNS Query Listen Enable

Use this parameter to display information only about logical interfaces that have the DNS query listen value you specify.

[-allow-lb-migrate {true|false}] - (DEPRECATED)-Load Balancing Migrate Allowed (privilege: advanced)

> This parameter has been deprecated and may be removed in a future version of ONTAP.

Use this parameter to display information only about logical interfaces for which load balancing migration is activated (true) or not activated (false).

[-lb-weight {load|0..100}] - Load Balanced Weight (privilege: advanced)

Use this parameter to display information only about logical interfaces that have the load balancing weight you specify.

[-failover-group <failover-group>] - Failover Group Name

Use this parameter to display information only about logical interfaces that are in the failover group you specify. Logical interfaces in the same failover group are capable of failing over to the same set of ports.

[-wwpn <text>] - FCP WWPN

Use this parameter to display information only about logical interfaces that have the Fibre Channel Protocol port identifier (World Wide Port Name) you specify.

[-address-family {ipv4|ipv6|ipv6z}] - Address family

Use this parameter to view the address family that is in use on the interface. Only IPv4 and IPv6 non-zoned addresses can be configured. Configuration of IPv6z addresses is not allowed.

[-comment <text>] - Comment

Use this parameter to display information only about logical interfaces that have the comment you specify.

[-ipspace <IPspace>] - IPspace of LIF

Use this parameter to display information only about logical interfaces on the IPspace you specify.

[-is-dns-update-enabled {true|false}] - Is Dynamic DNS Update Enabled?

Use this parameter to display information only about logical interfaces that have (true) or do not have (false) dynamic DNS updates enabled for them.

[-probe-port <integer>] - Probe-port for Cloud Load Balancer

Use this parameter display the probe-port for the logical interface in the Azure environment.

[-broadcast-domain <text>] - Broadcast Domain

Use this parameter to display the broadcast domain that contains the home port of the logical interface.

[-vserver-type <vserver type>] - Vserver Type

Use this parameter to display information only about logical interfaces owned by Vservers of the specified type.

[-rdma-protocols <roce>,…] - Required RDMA offload protocols

Use this parameter to display the logical interfaces associated with one or more RDMA protocols.

Examples

The following example displays general information about all logical interfaces.

```
cluster1::> network interface show
            Logical    Status     Network            Current       Current
Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port
Home
----------- ---------- ---------- ------------------ ------------- -------
----
cluster1
            cluster_mgmt
                       up/up      192.0.2.1/192      node0         e0M
true
            node0_mgmt1
                       up/up      192.0.2.2/192      node0         e0M
true
            node1_mgmt1
                       up/up      192.0.2.3/192      node1         e0M
true
Cluster
            node0_clus1
                       up/up      192.0.2.66/192     node0         e0a
true
            node0_clus2
                       up/up      192.0.2.67/192     node0         e0b
true
            node1_clus1
                       up/up      192.0.2.68/192     node1         e0a
true
            node1_clus2
                       up/up      192.0.2.69/192     node1         e0b
true
```

The following example displays failover information about all logical interfaces.

```
cluster1::> network interface show -failover
         Logical         Home                  Failover        Failover
Vserver  Interface       Node:Port             Policy          Group
-------- --------------- --------------------- ---------------
---------------
cluster1
         cluster_mgmt    node0:e0M             broadcast-domain-wide
                                                                Default
                         Failover Targets: node0:e0M,
                                           node0:e0d,
                                           node0:e0e,
                                           node0:e0f,
                                           node1:e0M,
                                           node1:e0d,
                                           node1:e0e,
                                           node1:e0f
         node0_mgmt1     node0:e0M             local-only      Default
                         Failover Targets: node0:e0M,
                                           node0:e0d,
                                           node0:e0e,
                                           node0:e0f
         node1_mgmt1     node1:e0M             local-only      Default
                         Failover Targets: node1:e0M,
                                           node1:e0d,
                                           node1:e0e,
                                           node1:e0f
Cluster
         node0_clus1     node0:e0a             local-only      Cluster
                         Failover Targets: node0:e0a,
                                           node0:e0b
         node0_clus2     node0:e0a             local-only      Cluster
                         Failover Targets: node0:e0b,
                                           node0:e0a
         node1_clus1     node1:e0a             local-only      Cluster
                         Failover Targets: node1:e0a,
                                           node1:e0b
         node1_clus2     node1:e0a             local-only      Cluster
                         Failover Targets: node1:e0b,
                                           node1:e0a
```

network interface start-cluster-check

(DEPRECATED)-Use the network interface check cluster-connectivity start command

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

Description

> This command has been deprecated and may be removed in a future version of ONTAP. Use the `network interface check cluster-connectivity start` command instead.

The `network interface start-cluster-check` command initiates an accessibility check from every logical interface to every aggregate. Automatic checks run periodically, but this command manually initiates a check immediately.

This command produces no direct output. Any errors encountered during the check are reported in the event log. See the event log show command for more information.

Examples

This example shows an execution of this command, with all parameters and output.

```
cluster1::> network interface start-cluster-check
```

Related Links

- event log show

network interface capacity show

Display the number of IP data LIFs capable of being configured on the cluster.

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

Description

The `network interface capacity show` command displays the number of IP LIFs of role *`data`* supported on the cluster, as well as the number of IP LIFs of role *`data`* currently configured on the cluster.

> The number of IP LIFs of role *`data`* that are supported on a node depends on the hardware platform and the Cluster's ONTAP version. If one or more nodes in the cluster cannot support additional LIFs, then none of the nodes in the cluster can support additional LIFs.

Examples

The following displays the IP data LIF capacity.

```
cluster1::> network interface capacity show
      IP Data LIF      IP Data LIF
  Supported Limit            Count
-----------------  ---------------
             1024              256
```

network interface capacity details show

Display details about the IP data LIFs capable of being configured on each node.

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

Description

The `network interface capacity details show` command displays the number of IP LIFs of role *`data`* that can be configured on each node, the number of IP data LIFs of role *`data`* that are supported on each node, and the number of IP data LIFs of role *`data`* that are configured to be homed on each node.

> The number of IP LIFs of role *`data`* that are supported on a node depends on the hardware platform and the Cluster's ONTAP version. If one or more nodes in the cluster cannot support additional LIFs, then none of the nodes in the cluster can support additional LIFs.

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>|local}]` - Node Name (privilege: advanced)

Use this parameter to specify the node for which to obtain data LIF capacity.

`[-capacity-for-node <integer>]` - Number of IP data LIFs that can be configured on the node (privilege: advanced)

This parameter specifies the number of IP LIFs of role *`data`* that can be configured on the node at the currently running ONTAP version. To view the version of a node, use the cluster image show command.

`[-limit-for-node <integer>]` - Number of IP data LIFs that are supported on the node (privilege: advanced)

This parameter specifies the number of IP LIFs of role *`data`* that are supported on the node at the current effective cluster version (ECV). To view the version of a node, use the cluster image show command.

`[-count-for-node <integer>]` - Number of IP data LIFs that are assigned to the node (privilege: advanced)

This parameter specifies the number of IP LIFs of role *`data`* currently configured to be homed on the node. To view LIFs homed on this node, use the network interface show -home-node command.

Examples

The following displays the IP data LIF capacity.

```
cluster1::> network interface capacity details show
                       IP Data LIF       IP Data LIF      IP Data LIF
             Node        Capacity   Supported Limit            Count
-----------------  ---------------  ---------------  ---------------
            node1              512              512              128
            node2              512              512              128
```

Related Links

- cluster image show
- network interface show

network interface check cluster-connectivity show

Display the Cluster Connectivity Log

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

Description

The `network interface check cluster-connectivity show` command displays the cluster check result log.

Parameters

`[-fields <fieldname>,…]`

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| `[-n ]` (privilege: advanced)

Selects IP addresses instead of LIF names.

| `[-instance ]` (privilege: advanced) }

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

`[-node {<nodename>|local}]` - Node (privilege: advanced)

Selects cluster health records that match the specified node name.

`[-record-number <integer>]` - Record Number (privilege: advanced)

Selects the cluster health records that match specified record number.

`[-date-time {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm]}]` - Record Date (privilege: advanced)

Selects the cluster health records that match the specified date and time.

`[-src-name <text>]` - Source LIF Name (privilege: advanced)

Selects the cluster health records that match the specified source LIF name.

[-src-addr <IP Address>] - Source IP Address (privilege: advanced)

Selects the cluster health records that match the specified source IP address.

[-dst-name <text>] - Destination LIF Name (privilege: advanced)

Selects the cluster health records that match the specified destination LIF name.

[-dst-addr <IP Address>] - Destination IP Address (privilege: advanced)

Selects the cluster health records that match the specified destination IP address.

[-loss {None|Partial|All}] - Packet Loss (privilege: advanced)

Selects the cluster health records that match the specified loss value.

Examples

The following example displays the cluster connectivity records.

```
cluster1::> network interface check cluster-connectivity show
                                  Source          Destination     Packet
Node   Date                       LIF             LIF             Loss
------ -------------------------- --------------- ---------------
-----------
node1
       8/26/2020 10:35:36 -04:00  node1_clus_1    node1_clus_2    none
       8/26/2020 10:35:37 -04:00  node1_clus_1    node2_clus_1    none
       8/26/2020 10:35:38 -04:00  node1_clus_1    node2_clus_2    none
       8/26/2020 10:35:39 -04:00  node1_clus_2    node1_clus_1    none
       8/26/2020 10:35:40 -04:00  node1_clus_2    node2_clus_1    none
       8/26/2020 10:35:41 -04:00  node1_clus_2    node2_clus_2    none
node2
       8/26/2020 10:35:36 -04:00  node2_clus_1    node2_clus_2    none
       8/26/2020 10:35:37 -04:00  node2_clus_1    node1_clus_1    none
       8/26/2020 10:35:38 -04:00  node2_clus_1    node1_clus_2    none
       8/26/2020 10:35:39 -04:00  node2_clus_2    node2_clus_1    none
       8/26/2020 10:35:40 -04:00  node2_clus_2    node1_clus_1    none
       8/26/2020 10:35:41 -04:00  node2_clus_2    node1_clus_2    none
12 entries were displayed.
```

network interface check cluster-connectivity start

Start the cluster check function

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

Description

The `network interface check cluster-connectivity start` command initiates an accessibility check from every logical interface to every aggregate. Automatic checks run periodically, but this command

manually initiates a check immediately.

This command produces no direct output. Any errors encountered during the check are reported via the cluster-check show and the event log. See the network interface check cluster-connectivity show and event log show commands for more information.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

This optionally selects the node on which to initiate the accessibility check. If no node is specified, the accessibility check is initiated on the local node.

Examples

This example shows an execution of this command with no parameter which will initiate the accessibility check on the local node.

```
cluster1::> network interface check cluster-connecivity start
```

Related Links

- network interface check cluster-connectivity show
- event log show

network interface check failover show

Discover if any LIFs might become inaccessible during a node outage, due to over-provisioning

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

Description

This command identifies logical interfaces (LIFs) at risk of becoming inaccessible if their hosting nodes were to experience an outage. The source-nodes parameter is the only required input.

The tuple <destination-nodes, vserver-name, lif-name> is sufficient to uniquely identify a record in the returned listing. All fields other than source-nodes can be filtered on in the usual fashion. There are some examples of this filtering below.

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.

[-destination-nodes <nodename>,…] - Set Of Nodes Over Capacity

Use this parameter to display the nodes an at-risk LIF or LIFs could fail over to.

[-vserver-name <vserver>] - Vserver Name

Use this parameter to display only LIFs on the Vserver you specify.

[-lif-name <lif-name>] - LIF Name

Use this parameter to display at-risk information only about the LIF or LIFs whose name you specify.

-source-nodes <nodename>,… - Nodes Going Down

List of nodes to test. At-risk LIFs currently hosted on these nodes will be identified. The list should contain no more than half the nodes in the cluster.

[-over-amount <integer>] - Amount Capacity Exceeded

Use this parameter to select only at-risk LIFs associated with a set of destination nodes whose amount over capacity matches the number you specify.

Note that the number of LIFs considered to be at risk may be higher than the actual amount over capacity a given set of nodes is. Once a given set of nodes is determined to be potentially over capacity, all LIFs whose set of failover target nodes is an exact match are marked as at risk. The amount over capacity is an upper bound on the number LIFs which could become unhosted if LIFs were to fail over in a random order, each to a target randomly selected from that LIF's configured failover targets.

[-failover-group <failover-group>] - Failover Group Name

Use this parameter to display information only about at-risk LIFs whose failover-group you specify.

[-failover-policy {system-defined|local-only|sfo-partner-only|disabled|broadcast-domain-wide}] - Failover Policy

Use this parameter to display information only about at-risk LIFs whose failover-policy you specify.

Examples

The following example shows all the at-risk LIFs for a specific two-node outage in a six-node cluster.

```
cluster1::> network interface check failover show -source-nodes
node1,node5

Destination Nodes: node2, node3, node4, node6
Amount Over Capacity: 2
Vserver             Logical Interface    Failover Group   Failover Policy
------------------  -------------------  ----------------
---------------------
vs0                 data1                Default          broadcast-
domain-wide
                    data2                Default          broadcast-
domain-wide
                    data3                Default          broadcast-
domain-wide
vs1                 data1                Custom_Name      broadcast-
domain-wide

Destination Nodes: node2
Amount Over Capacity: 1
Vserver             Logical Interface    Failover Group   Failover Policy
------------------  -------------------  ----------------
---------------------
vs0                 data6                Default          sfo-partner-only
vs1                 data7                Default          sfo-partner-only
```

The following example shows the same two-node outage scenario, but now with some filtering applied to the results.

```
cluster1::> network interface check failover show -source-nodes
node1,node5 -destination-nodes node2,node3,node4,node6 -failover-group
Def*

Destination Nodes: node2, node3, node4, node6
Amount Over Capacity: 2
Vserver             Logical Interface    Failover Group   Failover Policy
------------------  -------------------  ----------------
---------------------
vs0                 data1                Default          broadcast-
domain-wide
                    data2                Default          broadcast-
domain-wide
                    data3                Default          broadcast-
domain-wide
```

network interface dns-lb-stats show

Show the DNS load-balancer stats for this node

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

Description

The `network interface dns-lb-stats show` command displays the statistics for DNS load-balancing lookups for the zones belonging to the specified Vserver. These statistics represent the data for the Vserver on the local node. The following counts can be seen in the statistics output:

- success-count : Number of successful lookups.
- authoritative-count : Number of authoritative answers sent.
- nonauthoritative-count : Number of non authoritative answers sent.
- rr-set-missing-count : Number of times the RR set was missing.
- domain-missing-count : Number of times the domain was not be found.
- failure-count : Number of failed lookups.
- dropped-count : Number of lookups dropped.

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.

`[-vserver <vserver>]` - Vserver (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified Vservers.

`[-zone <text>]` - DNS Zone (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified DNS zones.

`[-success-count <integer>]` - Successful Lookup Count (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified number of successful lookups.

`[-authoritative-count <integer>]` - Authoritative Answer Count (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified number of authoritative answers sent.

`[-nonauthoritative-count <integer>]` - Non Authoritative Answer Count (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified number of non-authoritative answers sent.

[-rr-set-missing-count <integer>] - RR Set Missing Count (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified number of times the RR set was missing.

[-domain-missing-count <integer>] - Name Missing Count (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified number of times the domain was not found.

[-failure-count <integer>] - Failed Lookup Count (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified number of failed lookups.

[-dropped-count <integer>] - Dropped Count (privilege: advanced)

Use this parameter to display DNS load-balancer statistics only for the specified number of dropped lookups.

Examples

The following example displays stats for the zone "x.com".

```
cluster1::> network interface dns-lb-stats show -zone x.com
Vserver    DNS Zone           SUCCESS    AUTH  NOAUTH  NORR  NODOM  FAILED
DROP
---------  -----------------  -------   -----  ------ -----  -----  ------
-----
vs2
           x.com        5       5       0     0      0       0     0
```

network interface failover-groups add-targets

Add failover targets to a failover group

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

Description

The `network interface failover-groups add-targets` command enables you to add a list of failover targets such as network ports, interface groups, or VLANs to an existing logical interface failover group.

Parameters

-vserver <vserver> - Vserver Name

Use this parameter to specify the name of the Vservers from which this failover group is accessible.

-failover-group <text> - Failover Group Name

Use this parameter to specify the failover group that you want to extend.

-targets [node>:<port],… - Failover Targets

Use this parameter to specify the failover targets such as network ports, interface groups, or VLANs you want to add to the failover group.

Examples

This example shows the failover group "clyde" being extended to include additional failover targets.

```
cluster1::> network interface failover-group add-targets -vserver vs1
-failover-group clyde -targets xena1:e0c, xena1:e0d-100, xena2:a0a
```

network interface failover-groups create

Create a new failover group

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

Description

The `network interface failover-groups create` command creates a grouping of failover targets for logical interfaces on one or more nodes. Use this command to add a new network port or interface group to an existing failover group.

Parameters

-vserver <vserver> - Vserver Name

Use this parameter to specify the name of the Vservers from which this failover group is accessible.

-failover-group <text> - Failover Group Name

Use this parameter to specify the name of the logical interface failover group that you want to create.

-targets [node>:<port],… - Failover Targets

Use this parameter to specify the list of failover targets (network ports, interface groups, or VLANs on a node) belonging to this failover group.

Examples

The following example shows how to create a failover group named failover-group_2 containing ports e1e and e2e on node Xena.

```
cluster1::> network interface failover-groups create -vserver vs0
-failover-group failover-group_2 -targets xena:e1e,xena:e2e
```

network interface failover-groups delete

Delete a failover group

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

Description

The `network interface failover-groups delete` command deletes a logical interface failover group.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the name of the Vservers from which this failover group is accessible.

`-failover-group <text>` - Failover Group Name

Use this parameter to specify the name of the logical interface failover group to be deleted.

Examples

The following example shows how to delete a failover group named failover-group_2.

```
cluster1::> network interface failover-groups delete -vserver vs1
-failover-group failover-group_2
```

network interface failover-groups modify

Modify a failover group

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

Description

The `network interface failover-groups modify` command enables you modify the list of network ports, interface groups, or VLANs belonging to an existing logical interface failover group. The specified list will overwrite the existing list of network ports, interface groups, and VLANs currently belonging to the logical interface failover group.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the name of the Vserver(s) from which this failover group is accessible.

`-failover-group <text>` - Failover Group Name

Use this parameter to specify the failover group that you want to modify.

`[-targets [node>:<port],…]` - Failover Targets

Use this parameter to specify the network ports, interface groups, or VLANs you want to now belong to the failover group.

Examples

This example shows the failover group "clyde" being modified to now contain the specified network ports.

```
cluster1::> network interface failover-group modify -vserver vs1 -failover
-group clyde -targets xena1:e0c, xena1:e0d-100, xena2:a0a
```

network interface failover-groups remove-targets

Remove failover targets from a failover group

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

Description

The `network interface failover-groups remove-targets` command enables you to specify a list of failover targets such as network ports, interface groups, or VLANs to be removed from an existing logical interface failover group.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the name of the Vserver(s) from which this failover group is accessible.

`-failover-group <text>` - Failover Group Name

Use this parameter to specify the failover group that you want to remove failover targets from.

`-targets [node>:<port],…` - Failover Targets

Use this parameter to specify the failover targets such as network ports, interface groups, or VLANs you want to remove from the failover group.

Examples

This example shows the failover targets xena1:e0c and xena1:e0d-100 being removed from the failover group "clyde".

```
cluster1::> network interface failover-group remove-targets -vserver vs1
-failover-group clyde -targets xena1:e0c, xena1:e0d-100, xena2:a0a
```

network interface failover-groups rename

Rename a logical interface failover Group

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

Description

The `network interface failover-groups rename` command enables you to rename an existing logical interface failover group.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the name of the Vservers from which this failover group is accessible.

`-failover-group <text>` - Failover Group Name

Use this parameter to specify the failover group that you want to rename.

`-new-failover-group-name <text>` - New name

Use this parameter to specify the new name of the failover group.

Examples

This example shows the failover group "clusterwide" being renamed "clyde".

```
cluster1::> network interface failover-group rename -failover -vserver vs1
-failover-group clusterwide -new-failover-group-name clyde
```

network interface failover-groups show

Display logical interface failover groups

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

Description

The `network interface failover-groups show` command displays information about logical interface failover groups.

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.

`[-vserver <vserver>]` - Vserver Name

Use this parameter to display information only about the logical interface failover groups that have the target Vserver you specify.

`[-failover-group <text>]` - Failover Group Name

Use this parameter to display information only about the logical interface failover groups you specify.

`[-targets [node>:<port],…]` - Failover Targets

Use this parameter to display information only about the logical interface failover groups that have the failover target (physical port, interface group, or VLAN) you specify.

[-broadcast-domain <Broadcast Domain>] - Broadcast Domain

Use this parameter to display information only about the logical interface failover groups that have the broadcast domain you specify.

Examples

The following example displays information about all logical interface failover groups on a two node cluster.

```
cluster1::> network interface failover-groups show
                                   Failover
Vserver          Group             Targets
---------------- ----------------
--------------------------------------------
Cluster
                 Cluster
                                   node1:e1a, node1:e2a,
                                   node2:e1a, node2:e2a,
cluster1
                 Default
                                   node1:e0M, node1:e0a,
                                   node1:e0b, node1:e0c,
                                   node1:e0d, node2:e0M,
                                   node2:e0a, node2:e0b,
                                   node2:e0c, node2:e0d
```

network interface lif-weights show

Show the load-balancer LIF weights

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

Description

The `network interface lif-weights show` command displays the weights assigned to each LIF in a DNS load-balancing zone in a Vserver.

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.

[-vserver <vserver>] - Vserver (privilege: advanced)

Use this parameter to display information only for the specified Vservers.

[-zone <text>] - DNS Zone (privilege: advanced)

Use this parameter to display information only for the specified DNS zones.

[-address <IP Address>] - Network Address (privilege: advanced)

Use this parameter to display information only for the specified IP addresses.

[-weight <double>] - Load Balancer Weight (privilege: advanced)

Use this parameter to display information only for the specified load balancer weights

Examples

The following example displays LIF weights for vserver "vs1".

```
cluster1::> network interface lif-weights show -vserver vs1
                                Network
Vserver    DNS Zone             Address        Weight
---------  ------------------   ------------   -------
vs1
           a.com                4.4.4.4        12.4206
           x.com                1.1.1.1        12.4206
           x.com                10.72.46.236   12.4206
3 entries were displayed.
```

network interface service show

Display available interface services

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

Description

The `network interface service show ` command displays available services for IP LIFs and the TCP or UDP ports that each service listens on. The ports listed in this table correspond to well-known ports that each service can be expected to open a listening socket. Services that do not listen for ingress connections are presented with an empty port list.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| [-restrictions] (privilege: advanced)

The `network interface service show-restrictions` command displays available services for IP LIFs and usage restrictions for each service. The restrictions determine which LIFs are permitted to use each service and what restrictions the service implies for the LIFs that do use it.

| [-instance] }

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

[-service <LIF Service Name>] - Service Name

Selects services that match the specified service name.

[-ports <integer>,…] - TCP/UDP Port Numbers

Selects services that contain all IP ports in the specified list.

[-protocol-ports <text>,…] - Protocol: Port Numbers

Selects services that match the `<protocol>:<port>` combination.

[-vserver-policy <svc_vserver_policy>] - Vserver Restrictions

Selects services that match a specific vserver restriction.

[-failover-policy <svc_failover_policy>] - Failover Restrictions

Selects services that match a specific interface failover restriction.

Examples

The following example displays the built-in services.

```
cluster1::> network interface service show
Service                Protocol:Port
---------------------- -------------
intercluster-core      tcp:11104
                       tcp:11105
management-bgp         tcp:179

2 entries were displayed.
```

network interface service-policy add-service

Add an additional service entry to an existing service policy

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

Description

The `network interface service-policy add-service` command adds an additional service to an existing service-policy. When an allowed address list is specified, the list applies to only the service being added. Existing services included in this policy are not impacted.

Parameters

-vserver <vserver name> - Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver of the service policy to be updated.

`-policy <text>` - Policy Name (privilege: advanced)

Use this parameter to specify the name of service policy to be updated.

`-service <LIF Service Name>` - Service entry to be added (privilege: advanced)

Use this parameter to specify the name of service to be added to the existing service policy.

`[-allowed-addresses <IP Address/Mask>,…]` - Allowed address ranges for the service (privilege: advanced)

Use this parameter to specify a list of subnet masks for addresses that are allowed to access this service. Use the value `0.0.0.0/0` to represent the wildcard IPv4 address and `::/0` to represent the wildcard IPv6 address.

Examples

The following example shows the addition of a service to an existing service policy.

```
cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.

cluster1::> set advanced
Warning: These advanced commands are potentially dangerous; use them only
when directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y

cluster1::> network interface service-policy add-service -vserver cluster1
-policy default-intercluster -service management-ssh

cluster1::> set admin

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.
```

network interface service-policy clone

Clone an existing network service policy

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

Description

The `network interface service-policy clone` command creates a new service policy that includes the same services and allowed addresses as an existing policy. Once the new service policy has been created, it can be modified as necessary without impacting the original policy.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver of the service policy to be cloned.

`-policy <text>` - Policy Name (privilege: advanced)

Use this parameter to specify the name of the service policy to be cloned.

`-target-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the name of the vserver on which the new service policy should be created.

`-target-policy <text>` - New Service Policy Name (privilege: advanced)

Use this parameter to specify the name of the new service policy.

Examples

The following example shows the cloning of a service policy.

```
cluster1::> network interface service-policy show -vserver
cluster1,ipspace1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          custom1                    intercluster-core: 0.0.0.0/0
                                     management-core: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

ipspace1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
```

```
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

7 entries were displayed.

cluster1::> network interface service-policy clone -vserver cluster1
-policy custom1 -target-vserver ipspace1 -target-policy custom2

cluster1::> network interface service-policy show -vserver
cluster1,ipspace1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          custom1                    intercluster-core: 0.0.0.0/0
                                     management-core: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

ipspace1
          custom2                    intercluster-core: 0.0.0.0/0
                                     management-core: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

8 entries were displayed.
```

network interface service-policy create

Create a new service policy

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

Description

The `network interface service-policy create` command creates a new service policy with a list of included services. LIFs can reference this policy to control the list of services that they are able to transport on their network. Services can represent applications accessed by a LIF as well as applications served by this cluster.

Parameters

`-vserver <vserver name>` - Vserver

Use this parameter to specify the name of the Vserver on which the service policy will be created.

`-policy <text>` - Policy Name

Use this parameter to specify the name of service policy to be created.

`[-services <LIF Service Name>,…]` - Included Services

Use this parameter to specify a list of services that should be included in this policy.

`[-allowed-addresses <IP Address/Mask>,…]` - Allowed Addresses

Use this parameter to specify a list of subnet masks for addresses that are allowed to access the services in this policy. Use the value `0.0.0.0/0` to represent the wildcard IPv4 address and `::/0` to represent the wildcard IPv6 address.

Examples

The following example shows the creation of a service policy with no initial services.

```
cluster1::> network interface service-policy create -vserver cluster1
-policy empty

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0
empty                      -

4 entries were displayed.
```

The following example shows the creation of a new service policy with a specified service list.

```
cluster1::> network interface service-policy create -vserver cluster1
-policy custom -services intercluster-core,management-ssh

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- -------------------------
----------------------------------------
cluster1
          custom                     intercluster-core: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0
empty                      -

5 entries were displayed.
```

network interface service-policy delete

Delete an existing service policy

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

Description

The `network interface service-policy delete` command deletes an existing service policy.

Parameters

`-vserver <vserver name>` - Vserver

Use this parameter to specify the name of the Vserver of the service policy to be deleted.

`-policy <text>` - Policy Name

Use this parameter to specify the name of the service policy to be deleted.

Examples

The following example shows the deletion of a service policy.

```
cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          custom                     intercluster-core: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

4 entries were displayed.

cluster1::> network interface service-policy delete -vserver cluster1
-policy custom

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.
```

network interface service-policy modify-service

Modify a service entry in an existing service policy

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

Description

The `network interface service-policy modify-service` command modifies the policy-specific attributes of a service that is already included in a particular service policy. Other services in the policy are not

impacted by the change.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver of the service policy to be updated.

`-policy <text>` - Policy Name (privilege: advanced)

Use this parameter to specify the name of service policy to be updated.

`-service <LIF Service Name>` - Service entry to be modified (privilege: advanced)

Use this parameter to specify the name of service to be updated.

`-allowed-addresses <IP Address/Mask>,…` - Allowed address ranges for the service (privilege: advanced)

Use this parameter to specify a list of subnet masks for addresses that are allowed to access this service. Use the value `0.0.0.0/0` to represent the wildcard IPv4 address and `::/0` to represent the wildcard IPv6 address.

Examples

The following example shows the modification of a service on an existing service policy.

```
cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.

cluster1::> network interface service-policy modify-service -vserver
cluster1 -policy default-management -service management-ssh -allowed
-addresses 10.1.0.0/16

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 10.1.0.0/16
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.
```

network interface service-policy remove-service

Remove a service entry from an existing service policy

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

Description

The `network interface service-policy remove-service` command removes an individual service from an existing service policy. Other services in the the policy are not impacted by the change.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver of the service policy to be updated.

`-policy <text>` - Policy Name (privilege: advanced)

Use this parameter to specify the name of service policy to be updated.

`-service <LIF Service Name>` - Service entry to be removed (privilege: advanced)

Use this parameter to specify the name of service to be removed from the existing service policy.

Examples

The following example shows the removal of a service from an existing service policy.

```
cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- -------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.

cluster1::> network interface service-policy remove-service -vserver
cluster1 -policy default-management -service management-autosupport

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- -------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.
```

network interface service-policy rename

Rename an existing network service policy

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

Description

The `network interface service-policy rename` command assigns a new name to an existing service policy without disrupting the LIFs using the policy.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver of the service policy to be renamed.

`-policy <text>` - Policy Name (privilege: advanced)

Use this parameter to specify the name of the service policy to be renamed.

`-new-name <text>` - New Service Policy Name (privilege: advanced)

Use this parameter to specify the new name for the service policy.

Examples

The following example shows the renaming of a service policy.

```
cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          custom                     intercluster-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

4 entries were displayed.

cluster1::> network interface service-policy rename -vserver cluster1
-policy custom -new-name system

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          system                     intercluster-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

4 entries were displayed.
```

network interface service-policy restore-defaults

Restore default settings to a service policy

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

Description

The `network interface service-policy restore-defaults` command restores a built-in service-policy to its original state. The default list of services replaces any customizations that have been applied by an administrator. All included services will be updated to use the default allowed address list.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

Use this parameter to specify the name of the Vserver of the service policy to be restored.

`-policy <text>` - Policy Name (privilege: advanced)

Use this parameter to specify the name of the service policy to be restored.

Examples

The following example shows the restoration of a service policy's default settings.

```
cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 10.1.0.0/16
                                     management-ssh: 10.1.0.0/16
                                     management-https: 10.1.0.0/16
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.

cluster1::> network interface service-policy restore-defaults -vserver
cluster1 -policy default-intercluster

cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.
```

network interface service-policy show

Display existing service policies

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

Description

The `network interface service-policy show` command displays existing service policies.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the fields that you specify.

| [-instance] }

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

[-vserver <vserver name>] - Vserver

Selects service policies that match the specified vserver name.

[-policy <text>] - Policy Name

Selects service policies that match the specified service policy name.

[-services <LIF Service Name>,…] - Included Services

Selects service policies that contain all services in the specified list of service names.

[-service-allowed-addresses <text>,…] - Service: Allowed Addresses

Selects service policies that contain all "`<service>:<allowed-addresses>`" in the specified list of addresses.

Examples

The following example displays the built-in service policies.

```
cluster1::> network interface service-policy show -vserver cluster1
Vserver   Policy                     Service: Allowed Addresses
--------- --------------------------
----------------------------------------
cluster1
          default-intercluster       intercluster-core: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-management         management-core: 0.0.0.0/0
                                     management-autosupport: 0.0.0.0/0
                                     management-ssh: 0.0.0.0/0
                                     management-https: 0.0.0.0/0
default-route-announce     management-bgp: 0.0.0.0/0

3 entries were displayed.
```

network ipspace commands

network ipspace create

Create a new IPspace

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

Description

IPspaces are distinct IP address spaces in which Storage Virtual Machines (SVMs) reside. The "Cluster" IPspace and "Default" IPspace are created by default. You can create more custom IPspaces when you need your SVMs to have overlapping IP addresses, or you need more control over networking configurations for cluster peering. Reference the "Network Management Guide" for the limit of how many custom IPspaces are supported on your system..

Parameters

`-ipspace <IPspace>` - IPspace name

The name of the IPspace to be created.

- The name must contain only the following characters: A-Z, a-z, 0-9, ".", "-" or "_".
- The first character of each label, delimited by ".", must be one of the following characters: A-Z or a-z.
- The last character of each label, delimited by ".", must be one of the following characters: A-Z, a-z or 0-9.
- The maximum supported length is 47 characters.
- The system reserves the following names: "all", "local" and "localhost".
- The system provides the following IPspaces: "Cluster" and "Default".

Examples

The following example creates IPspace "ips1".

```
cluster1::> network ipspace create -name ips1
```

network ipspace delete

Delete an IPspace

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

Description

Delete an IPspace that contains no ports or Vservers.

Parameters

`-ipspace <IPspace>` - IPspace name

The name of the IPspace to be deleted. If the IPspace is associated with one or more logical-interfaces, you must delete them before you can delete the IPspace.

Examples

The following example deletes the IPspace "ips1".

```
cluster1::> network ipspace delete -ipspace ips1
```

network ipspace rename

Rename an IPspace

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

Description

Rename an IPspace.

Parameters

`-ipspace <IPspace>` - IPspace name

The name of the IPspace to be renamed.

`-new-name <IPspace>` - New Name

The new name for the IPspace.

- The name must contain only the following characters: A-Z, a-z, 0-9, ".", "-" or "_".
- The first character of each label, delimited by ".", must be one of the following characters: A-Z or a-z.
- The last character of each label, delimited by ".", must be one of the following characters: A-Z, a-z or 0-9.
- The maximum supported length is 47 characters.
- The system reserves the following names: "all", "cluster", "local" and "localhost".

Examples

The following example renames IPspace "ips1" to "ips2".

```
cluster1::> network ipspace rename -ipspace ips1 -new-name ips2
```

network ipspace show

Display IPspace information

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

Description

Display network IPspaces.

Parameters

{ [-fields <fieldname>,…]

Specify the fields to be displayed for each IPspace.

| [-instance] }

Display all parameters of the IPspace objects.

[-ipspace <IPspace>] - IPspace name

Display the names of the IPspaces.

[-ports [node>:<port],…] - Ports

The list of network ports assigned to each IPspace.

[-broadcast-domains <Broadcast Domain>,…] - Broadcast Domains

The list of broadcast domains that belong to the IPspace.

[-vservers <vserver name>,…] - Vservers

The list of Vservers assigned to each IPspace.

Examples

The following example displays general information about IPspaces.

```
cluster1::> network ipspace show
IPspace             Vserver List                  Broadcast Domains
------------------- ----------------------------
----------------------------
Cluster
                    Cluster                       -
Default
                    cluster1, vs1, vs2            br1, br2, br3
2 entries were displayed.
```

network ndp commands

network ndp default-router delete-all

Delete default routers on a given IPspace

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

Description

The `network ndp default-router delete-all` command deletes default router lists from the specified IPspace.

Parameters

`-ipspace <IPspace>` - IPspace Name (privilege: advanced)

Use this parameter to specify the IPspace where the default routers are to be deleted.

Examples

The following example deletes default routers from IPspace ips1.

```
cluster1::*> network ndp default-router delete-all -ipspace ips1
```

network ndp default-router show

Display default routers

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

Description

The `network ndp default-router show` command displays Neighbor Discovery Protocol (NDP) default routers learned on a specified port.

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>|local}]` - Node (privilege: advanced)

Displays the NDP default routers from the specified node.

`[-ipspace <IPspace>]` - IPspace name (privilege: advanced)

Displays the NDP default routers from the specified IPspace.

`[-port {<netport>|<ifgrp>}]` - Port (privilege: advanced)

Displays the NDP default routers from the specified port.

`[-router-addr <IP Address>]` - Router Address (privilege: advanced)

Displays the default routers that have the specified IPv6 addresses.

`[-flag {none|managed-address-DHCPv6|other-DHCPv6}]` - Flag (privilege: advanced)

Displays the default routers that have the specified flag. The flag indicates whether addresses are available via DHCPv6 or other configuration information is available via DHCPv6.

[-expire-time {[<integer>d][<integer>h][<integer>m][<integer>s]|never|expired}] - Expire Time (privilege: advanced)

Displays the default routers that have the specified expire time.

Examples

The following example displays NDP default routers on local port e0f.

```
cluster1::*> network ndp default-router show -port e0f -node local

Node: node1
IPspace: Default
Port     Router Address             Flag            Expire Time
-------- -------------------------- --------------  --------------
e0f      fe80::5:73ff:fea0:107      none            0d0h23m9s
```

network ndp neighbor create

Create a static NDP neighbor entry

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

Description

The `network ndp neighbor create` command creates a static Neighbor Discovery Protocol (NDP) neighbor entry within a Vserver.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver on which the NDP neighbor is to be created.

-neighbor <IP Address> - Neighbor Address (privilege: advanced)

Use this parameter to specify the neighbor's IPv6 address.

-mac-address <MAC Address> - MAC Address (privilege: advanced)

Use this parameter to specify the neighbor's MAC address.

Examples

The following example creates a NDP neighbor entry within Vserver vs0.

```
cluster1::*> network ndp neighbor create -vserver vs0 -neighbor 20:20::20
-mac-address 10:10:10:0:0:1
```

network ndp neighbor delete

Delete a static NDP neighbor entry

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

Description

The `network ndp neighbor delete` command deletes a static Neighbor Discovery Protocol (NDP) neighbor from a Vserver.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver on which the NDP neighbor is to be deleted.

`-neighbor <IP Address>` - Neighbor Address (privilege: advanced)

Use this parameter to specify the neighbor's IPv6 address.

Examples

The following example deletes a NDP neighbor entry within Vserver vs0.

```
cluster1::*> network ndp neighbor delete -vserver vs0 -neighbor 20:20::20
```

network ndp neighbor show

Display static NDP neighbor entries

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

Description

The `network ndp neighbor show` command displays a group of static Neighbor Discovery Protocol (NDP) neighbors within one or more Vservers. You can view static NDP neighbors within specified Vservers, neighbors with specified IPv6 address, and neighbors with specified MAC address.

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.

`[-vserver <vserver name>]` - Vserver Name (privilege: advanced)

Displays the static NDP neighbors that have the specified Vserver as their origin.

[-neighbor <IP Address>] - Neighbor Address (privilege: advanced)

Displays the static NDP neighbors that have the specified IPv6 address.

[-mac-address <MAC Address>] - MAC Address (privilege: advanced)

Displays the static NDP neighbors that have the specified MAC address.

Examples

The following example displays all of the static NDP neighbors configured on Vserver vs0.

```
cluster1::*> network ndp neighbor show -vserver vs0
Vserver            Neighbor                 MAC Address
------------------ ------------------------ -------------------
vs0
                   10:10::10                04:04:04:04:04:04
                   20:20::20                01:01:01:01:01:01
2 entries were displayed.
```

network ndp neighbor active-entry delete

Delete active neighbor entry from a System or Admin Vserver

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

Description

The `network ndp neigbhor active-entry delete` command deletes a Network Discovery Protocol (NDP) neighbor entry on the specified port from a given Vserver's subnet group.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node on which the neighbor entry is to be deleted.

-vserver <vserver> - System or Admin Vserver Name (privilege: advanced)

Use this parameter to specify the System or Admin Vserver on which the neighbor entry is to be deleted.

-subnet-group <IP Address/Mask> - Subnet Group (privilege: advanced)

Use this parameter to specify the subnet group from which the neighbor entry is to be deleted.

-neighbor <IP Address> - Neighbor (privilege: advanced)

Use this parameter to specify the IPv6 address of the neighbor entry which is to be deleted.

-port {<netport>|<ifgrp>} - Port (privilege: advanced)

Use this parameter to specify the port on which the neighbor entry is to be deleted.

Examples

The following example deletes a neighbor entry from the Admin Vserver cluster1:

```
cluster1::*> network ndp neighbor active-entry delete -vserver cluster1
-node local -subnet-group ::/0 -neighbor fe80:4::5:73ff:fea0:107 -port e0d
```

network ndp neighbor active-entry show

Display active neighbor entries organized by Vserver

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

Description

The `network ndp neighbor active-entry show` command displays Network Discovery Protocol (NDP) neighbor cache entries on one or more nodes. You can view ndp neighbors within specified nodes and within specified System or Admin Vservers.

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.

| `[-verbose ]` (privilege: advanced)

Displays the expire time, state, is-router, and probe count fields.

| `[-instance ]` (privilege: advanced) }

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

`[-node {<nodename>|local}]` - Node (privilege: advanced)

Displays the NDP neighbors from the specified node.

`[-vserver <vserver>]` - System or Admin Vserver Name (privilege: advanced)

Displays the NDP neighbors from the specified System or Admin Vserver. Data and Node Vservers do not have dynamic NDP neighbors.

`[-subnet-group <IP Address/Mask>]` - Subnet Group (privilege: advanced)

Displays the NDP neighbors in the specified subnet group.

`[-neighbor <IP Address>]` - Neighbor (privilege: advanced)

Displays the NDP neighbors that have the specified IPv6 address.

`[-port {<netport>|<ifgrp>}]` - Port (privilege: advanced)

Displays the NDP neighbors on the specified port.

[-mac-address <MAC Address>] - MAC Address (privilege: advanced)

Displays the NDP neighbors have the specified MAC address.

[-expire-time {[<integer>d][<integer>h][<integer>m][<integer>s]|never|expired}] - Expire Time (privilege: advanced)

Displays the NDP neighbors have the specified expire time.

[-state {<nostate|incomplete|reachable|stale|delay|probe|unknown>}] - State (privilege: advanced)

Displays the NDP neighbors in the specified state.

[-is-router {true|false}] - Is Router (privilege: advanced)

Displays the NDP neighbor which is a router.

[-probe-count <integer>] - Probe Count (privilege: advanced)

Displays the NDP neighbors with the specified probe count. Probe count is the number of times that this neighbor's MAC address has been queried.

[-is-static {true|false}] - Is Static (privilege: advanced)

Displays the NDP neighbors which are statically configured.

Examples

The following example displays NDP neighbors on the Admin Vserver cluster1:

```
cluster1::*> network ndp neighbor active-entry show -vserver cluster1

Node: node1
Vserver: cluster1
Subnet Group: ::/0
Neighbor                       MAC Address          Port
------------------------------ -------------------- --------
fe80:4::5:73ff:fea0:107        00:05:73:a0:01:07    e0d
fe80:4::226:98ff:fe0c:b6c1     00:26:98:0c:b6:c1    e0d
fe80:4::4255:39ff:fe25:27c1    40:55:39:25:27c1     e0d
3 entries were displayed.
```

network ndp prefix delete-all

Delete IPv6 prefixes on a given IPspace

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

Description

The `network ndp prefix delete-all` command deletes all prefixes learned from the specified IPspace.

Parameters

`-ipspace <IPspace>` - IPspace Name (privilege: advanced)

Use this parameter to specify the IPspace where the IPv6 prefixes are to be deleted.

Examples

The following example deletes all IPv6 prefixes within IPspace ips1.

```
cluster1::*> network ndp prefix delete-all -ipspace ips1
```

network ndp prefix show

Display IPv6 prefixes

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

Description

The `network ndp prefix show` command displays IPv6 prefixes on one or more nodes.

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.

| `[-verbose ]` (privilege: advanced)

Displays the valid-lifetime, preferred-lifetime, origin and advertising-router fields.

| `[-instance ]` (privilege: advanced) }

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

`[-node {<nodename>|local}]` - Node (privilege: advanced)

Displays the IPv6 prefixes from the specified node.

`[-ipspace <IPspace>]` - IPspace name (privilege: advanced)

Displays the IPv6 prefixes from the specified IPspace.

`[-port {<netport>|<ifgrp>}]` - Port (privilege: advanced)

Displays the IPv6 prefixes on the specified port.

`[-prefix <IP Address/Mask>]` - Prefix (privilege: advanced)

Displays the IPv6 prefixes with the specified prefix value.

`[-flag {none|on-link|autonomous|on-link-autonomous}]` - Flag (privilege: advanced)

Displays the IPv6 prefixes with the specified flag. The flag indicates whether a prefix is on-link and whether it can be used in autonomous address configuration.

[-valid-lifetime {<unsigned integer>|infinity}] - Valid Lifetime (privilege: advanced)

Displays the IPv6 prefixes having the specified valid lifetime in seconds.

[-preferred-lifetime {<unsigned integer>|infinity}] - Preferred Lifetime (privilege: advanced)

Displays the IPv6 prefixes having the specified preferred lifetime in seconds.

[-expire-time {[<integer>d][<integer>h][<integer>m][<integer>s]|never|expired}] - Expire Time (privilege: advanced)

Displays the IPv6 prefixes having the specified expire time.

[-origin {router-advertise|renumber-request|static|kernel|unknown}] - Origin of the Prefix (privilege: advanced)

Displays the IPv6 prefixes with the specified origin.

[-advertising-router <IP Address>,…] - Router that Advertised the Prefix (privilege: advanced)

Displays the IPv6 prefixes which are propagated by the specified router list.

Examples

The following example displays IPv6 prefixes on port e0f.

```
cluster1::*> network ndp prefix show -port e0f -node local

Node: node1
IPspace: Default
Port      Prefix                    Flag               Expire Time
--------- ------------------------- ------------------ -------------
e0f       fd20:8b1e:b255:814e::/64  on-link-autonomous 29d23h56m48s
```

network options commands

network options cluster-health-notifications modify

cluster health notification options

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

Description

This command enables or disables cluster health notifications on the specified node.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node for which the cluster health notification status will be modified.

[-enabled {true|false}] - Cluster Health Notifications Enabled

Setting this parameter to *true* enables cluster health notification. Setting it to *false* disables cluster health notification.

Examples

The following example modifies the cluster health notification status for a node:

```
cluster1::> network options cluster-health-notifications modify -node
node1 -enabled true
```

network options cluster-health-notifications show

Display cluster health notification options

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

Description

The `network options cluster-health-notifications show` command displays whether the node's cluster health notifications are enabled.

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>|local}] - Node

This parameter specifies the node for which the cluster health notification status will be displayed.

[-enabled {true|false}] - Cluster Health Notifications Enabled

Selects the entries that match this parameter value.

Examples

The following example displays the cluster health notification status for a node:

```
cluster1::> network options cluster-health-notifications show -node node1
                                  Node: node1
 Cluster Health Notifications Enabled: true
```

network options detect-switchless-cluster modify

Modify the status of switchless cluster detection

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

Description

This command enables or disables the automatic detection of a switchless cluster. A switchless cluster consists of two nodes where the cluster ports are directly connected without a switch between them.

Parameters

`[-enabled {true|false}]` - Enable Switchless Cluster Detection (privilege: advanced)

This parameter specifies whether switchless cluster detection is enabled or not. Setting this parameter to *`true`* enables switchless cluster detection.

Examples

```
The following example enables switchless cluster detection:
        cluster1::*> network options detect-switchless-cluster modify
-enabled true
```

network options detect-switchless-cluster show

Display the status of switchless cluster detection

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

Description

The network options detect-switchless-cluster show command displays whether switchless cluster detection is enabled.

Examples

```
The following example displays whether switchless cluster detection is
enabled:
      cluster1::*> network options detect-switchless-cluster show
      Enable Detect Switchless Cluster: true
```

network options ipv6 modify

Modify IPv6 options

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

Description

This command sets the state of IPv6 options for the cluster.

Parameters

`[-enabled <true>]` - IPv6 Enabled

Setting this parameter to *`true`* enables IPv6 for the cluster. IPv6 cannot be disabled once it is enabled for the cluster. Call technical support for guidance regarding disabling IPv6.

`[-is-ra-processing-enabled {true|false}]` - Router Advertisement (RA) Processing Enabled

Setting this parameter to *`true`* enables cluster to process IPv6 router advertisements. Setting it to *`false`* disables router advertisement processing by the cluster.

Examples

```
The following example enables IPv6 for the cluster:
    cluster1::> network options ipv6 modify -enabled true

The following example enables IPv6 Router Advertisement processing for the
cluster:
    cluster1::> network options ipv6 modify -is-ra-processing-enabled true

The following example disables IPv6 Router Advertisement processing for
the
cluster:
    cluster1::> network options ipv6 modify -is-ra-processing-enabled
false
```

network options ipv6 show

Display IPv6 options

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

Description

This command displays the current state of IPv6 options for the cluster.

Examples

```
cluster1::> network options ipv6 show

                                IPv6 Enabled: false
Router Advertisement (RA) Processing Enabled: false
```

network options load-balancing modify

Modify load balancing algorithm

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

Description

This command sets the state of geometric mean algorithm for load balancing

Parameters

`[-enable {true|false}]` - Geometric Mean Algorithm for Load Balancing (privilege: advanced)

Setting this parameter to *`true`* enables the geometric mean algorithm for load balancing. Setting it to *`false`* disables the geometric mean algorithm for the cluster.

Examples

```
The following example will enable the geometric mean algorithm for load
      balancing.
      cluster1::> network options load-balancing modify -enable true
The following example will disable the geometric mean algorithm for load
      balancing.
      cluster1::> network options load-balancing modify -enable false
```

network options load-balancing show

Display load balancing algorithm

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

Description

This command displays the use of geometric mean load balancing algorithm.

Examples

```
cluster1::> network options load-balancing show
Geometric Mean Algorithm for load balancing: false
```

network options multipath-routing modify

Modify multipath-routing

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

Description

The `network options multipath-routing modify` command is used to modify cluster-wide multipath routing configuration.

Parameters

`[-is-enabled {true|false}]` - Is Multipath Routing Enabled (privilege: advanced)

This parameter specifies whether multipath routing configuration is enabled or not. Setting this parameter to _ `true`_ enables multipath routing for all nodes in the cluster.

Examples

```
cluster1::> network options multipath-routing modify -is-enabled true
```

network options multipath-routing show

Display multipath-routing

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

Description

The `network options multipath-routing show` command displays the multipath routing configuration for the cluster.

Examples

```
cluster1::> network options multipath-routing show
                Is Multipath Routing Enabled: false
```

network options port-health-monitor disable-monitors

Disable one or more port health monitors

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

Description

This command disables the given port health monitors for the given IPspaces in the cluster.

Parameters

`-ipspace <IPspace>` - IPspace Name (privilege: advanced)

The name of the IPspace for which the specified port health monitors are disabled.

`-health-monitors {l2-reachability|link-flapping|crc-errors|vswitch-link}` - List of Port Health Monitors to Disable (privilege: advanced)

The port health monitors to disable.

Examples

The following example disables the "l2_reachability" health monitor for the "Default" IPspace.

> The status of the "link_flapping" monitor is unaffected by the command.

```
cluster1::*> network options port-health-monitor show

IPspace         Enabled Port Health Monitors
-------------   ----------------------------
Cluster         l2_reachability,
                link_flapping
Default         l2_reachability,
                link_flapping
2 entries were displayed.

cluster1::*> network options port-health-monitor disableMonitors -ipspace
Default -health-monitors l2_reachability

cluster1::*> network options port-health-monitor show

IPspace         Enabled Port Health Monitors
-------------   ----------------------------
Cluster         l2_reachability,
                link_flapping
Default         link_flapping
2 entries were displayed.
```

network options port-health-monitor enable-monitors

Enable one or more port health monitors

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

Description

This command enables the given port health monitors for the given IPspaces in the cluster.

Parameters

`-ipspace <IPspace>` - IPspace Name (privilege: advanced)

The name of the IPspace for which the specified port health monitors are enabled.

`-health-monitors {l2-reachability|link-flapping|crc-errors|vswitch-link}` - List of Port Health Monitors to Enable (privilege: advanced)

The port health monitors to enable. Upon enabling the *`l2_reachability`* health monitor, it runs in an

"unpromoted" state. While in this state, the monitor does not mark any ports as unhealthy due to the l2_reachability health check. The monitor is promoted in the "Cluster" IPspace when the "Cluster" broadcast domain is found to have passed the *l2_reachability* health check. An EMS event called "vifmgr.hm.promoted" event is generated when the health monitor is promoted for the IPspace.

Examples

The following example enables the "l2_reachability" health monitor for the "Default" IPspace:

> The status of the "link_flapping" monitor is unaffected by the command.

```
cluster1::*> network options port-health-monitor show

IPspace         Enabled Port Health Monitors
-------------   ----------------------------
Cluster         l2_reachability,
                link_flapping
Default         link_flapping
2 entries were displayed.

cluster1::*> network options port-health-monitor enableMonitors -ipspace
Default -health-monitors l2_reachability

cluster1::*> network options port-health-monitor show

IPspace         Enabled Port Health Monitors
-------------   ----------------------------
Cluster         l2_reachability,
                link_flapping
Default         l2_reachability,
                link_flapping
2 entries were displayed.
```

network options port-health-monitor modify

Modify port health monitors configuration

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

Description

This command modifies the enabled port health monitors for the given IPspaces in the cluster.

Parameters

`-ipspace <IPspace>` - IPspace Name (privilege: advanced)

The name of the IPspace for which enabled port health monitors are modified.

`[-health-monitors {l2-reachability|link-flapping|crc-errors|vswitch-link}]` - List of Enabled Port Health Monitors (privilege: advanced)

All of the port health monitors that you want to enable. This command enables any port health monitors in this list that are currently disabled, and it disables any currently enabled monitors that are not in this list. Upon enabling the *`l2_reachability`* health monitor, it runs in an "unpromoted" state. While in this state, the monitor does not mark any ports as unhealthy due to the l2_reachability health check. The monitor is promoted in the "Cluster" IPspace when the "Cluster" broadcast domain is found to have passed the *`l2_reachability`* health check. An EMS event called "vifmgr.hm.promoted" event is generated when the health monitor is promoted for the IPspace.

Examples

The following example modifies the port health monitor configuration of the "Default" IPspace such that only the "link_flapping" port health monitor is enabled. enabled for all IPspaces in the cluster.

Only the specified monitor is enabled after the modify command is issued.

```
cluster1::*> network options port-health-monitor show

IPspace         Enabled Port Health Monitors
-------------   ----------------------------
Cluster         l2_reachability,
                link_flapping
Default         l2_reachability,
                link_flapping
2 entries were displayed.

cluster1::*> network options port-health-monitor modify -ipspace Default
-health-monitors link_flapping

cluster1::*> network options port-health-monitor show

IPspace         Enabled Port Health Monitors
-------------   ----------------------------
Cluster         l2_reachability,
                link_flapping
Default         link_flapping
2 entries were displayed.
```

network options port-health-monitor show

Display port health monitors configuration

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

Description

This command displays the enabled port health monitors for the IPspaces in the cluster.

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.

[-ipspace <IPspace>] - IPspace Name (privilege: advanced)

Displays the port health monitors that are enabled only for the given IPspace name.

[-health-monitors {l2-reachability|link-flapping|crc-errors|vswitch-link}] - List of Enabled Port Health Monitors (privilege: advanced)

Displays the IPspaces that have the given monitors enabled.

Examples

The following example lists all port health monitors that are enabled for all IPspaces in the cluster.

```
cluster1::*> network options port-health-monitor show

IPspace         Enabled Port Health Monitors
-------------   ----------------------------
Cluster         l2_reachability,
                link_flapping
Default         l2_reachability,
                link_flapping
2 entries were displayed.
```

network options send-soa modify

Modify Send SOA settings

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

Description

This command sets the status of sending statement of authority record in the DNS response.

Parameters

[-enable {true|false}] - Enable Sending SOA (privilege: advanced)

Setting this parameter to *true* enables sending the statement of authority (SOA) record in the DNS response. Setting it to *false* disables sending the statement of authority (SOA) record in the DNS

response for the cluster.

Examples

```
The following example will enable the sending of statement of authority
(SOA)
      in the DNS response.
      cluster1::> network options send-soa modify -enable true
The following example will disable the sending of statement of authority
(SOA)
      in the DNS response.
      cluster1::> network options send-soa modify -enable false
```

network options send-soa show

Display Send SOA settings

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

Description

This command displays whether sending the statement of authority record (SOA) in the DNS response is enabled or not.

Examples

```
cluster1::> network options send-soa show
Enable sending SOA: true
```

network options switchless-cluster modify

Modify switchless cluster network options

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

Description

This command sets whether the cluster network is in switchless or switched mode. A switchless cluster is physically formed by connecting two nodes back-to-back, without a switch between them.

Parameters

`[-enabled {true|false}]` - Enable Switchless Cluster (privilege: advanced)

This parameter specifies whether the switchless cluster is enabled or not. Setting this parameter to *`true`* enables the switchless cluster.

Examples

```
The following example enables the switchless cluster:
        cluster1::*> network options switchless-cluster modify -enabled
true
```

network options switchless-cluster show

Display switchless cluster network options

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

Description

The network options switchless-cluster show command displays the attributes of a switchless cluster.

Examples

```
The following example displays the attributes of the switchless cluster:
            cluster1::*> network options switchless-cluster show
            Enable Switchless Cluster: true
```

network port commands

network port delete

Delete a network port

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

Description

The `network port delete` command deletes a network port that is no longer physically present on the storage system.

Parameters

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

This specifies the node on which the port is located.

`-port {<netport>|<ifgrp>}` - Port

This specifies the port to delete.

Examples

The following example deletes port e0c from a node named node0. The command works only when the port does not physically exist on the storage system.

```
cluster1::*> network port delete -node node0 -port e0c
```

network port modify

Modify network port attributes

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

Description

The `network port modify` command enables you to change the maximum transmission unit (MTU) setting, autonegotiation setting, administrative duplex mode, and administrative speed of a specified network port.

The MTU of ports that belong to broadcast-domains must be updated through the `broadcast-domain modify` command.

Modification of a port's IPspace will only work before a node is added to a cluster, when the cluster version is below ONTAP 8.3, or when the node is offline. To change the IPspace of a port once the node is in a ONTAP 8.3 cluster, the port should be added to a broadcast-domain that belongs to that IPspace.

Parameters

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

Use this parameter to specify the node on which the port is located.

`-port {<netport>|<ifgrp>}` - Port

Use this parameter to specify the port that you want to modify.

`[-mtu <integer>]` - MTU

The port's MTU setting. The default setting for ports in the "Cluster" IPspace is 9000 bytes. All other ports use a default value of 1500 bytes.

`[-autonegotiate-admin {true|false}]` - Auto-Negotiation Administrative

Whether the port uses Ethernet autonegotiation to determine the highest speed and duplex mode that the port and its endpoint can support. The default setting when you create a port is `true` .

`[-duplex-admin {auto|half|full}]` - Duplex Mode Administrative

The administrative setting for the port's duplex mode. This is the duplex mode that you prefer the port to use. Depending on network limitations, the operational value can be different from the administrative setting. The default setting when you create a port is `full` .

`[-speed-admin {auto|10|100}]` - Speed Administrative

The administrative speed setting, in megabits per second. This is the speed setting that you prefer the port to use. Depending on network limitations, the operational value can be lower than the administrative setting.

`[-flowcontrol-admin {none|receive|send|full|pfc}]` - Flow Control Administrative

The administrative flow control setting of the port. this is the flow control setting that you prefer the port to use. Depending on network and port limitations, the operational value can be different from the administrative setting.

[-up-admin {true|false}] - Up Administrative (privilege: advanced)

The administrative state of the port. If set to `true` , the port is used if it is operational. If set to `false` , the port is configured down.

[-ipspace <IPspace>] - IPspace Name

Use this parameter to specify the IPspace the network port is assigned to. Modification of a port's IPspace will only work before a node is added to a cluster, when the cluster version is below ONTAP 8.3, or when the node is offline. To change the IPspace of a port once the node is in a ONTAP 8.3 cluster, the port should be added to a broadcast-domain that belongs to that IPspace. If there is an inconsistency between the broadcast-domain and IPspace, this parameter can be set to bring the IPspace into alignment with the broadcast-domain.

[-ignore-health-status {true|false}] - Ignore Port Health Status (privilege: advanced)

Use this parameter to specify that the system ignore network port health status of the specified port for the purpose of hosting a logical interface.

[-pfc-queues-admin <integer>,…] - List of PFC Queues Administrative

Use this parameter to set the list of PFC queues.

Examples

The following example modifies port e0a on a node named node0 not to use auto-negotiation, to preferably use half duplex mode, and to preferably run at 100 Mbps.

```
cluster1::> network port modify -node node0 -port e0a -autonegotiate-admin
false -duplex-admin half -speed-admin 100
```

network port show-address-filter-info

Print the port's address filter information

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

Description

The `network port show-address-filter-info` command displays information about the port's address filter.

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

Use this parameter to specify the node.

-port {<netport>|<ifgrp>} - Port

Use this parameter to specify the port. For example, e0c.

[-num-total <integer>] - Total Number Of Entries

Use this parameter to specify the total number of entries.

[-num-used <integer>] - Number Of Used Entries

Use this parameter to specify the number of used entries.

[-used-entries <text>,…] - The Used Entries

Use this parameter to list the used entries.

Examples

The following example displays information of the given port's address filter on the specified node of the cluster.

```
cluster1::*> network port show-address-filter-info -node local -port e0c

Node: node1
        Total Number       Number of
Port      of Address    Used Address     Used Address
Name  Filter Entries  Filter Entries       Filter Entries
----  --------------  --------------   ------------------
e0c             1328               3     U 0 a0 98 40 e 6
                                          M 1 80 c2 0 0 e
                                          M 1 0 5e 0 0 fb
```

network port show

Display network port attributes

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

Description

The `network port show` command displays information about network ports. The command output indicates any inactive links, and lists the reason for the inactive status.

Some parameters can have "administrative" and "operational" values. The administrative setting is the preferred value for that parameter, which is set when the port is created or modified. The operational value is the actual current value of that parameter. Administrative and operational settings are not shown for virtual ports, '-' will be displayed. See the physical port hosting the target virtual port for these values.

If the operational duplex mode and speed of a port cannot be determined (for instance, if the link is down), that port's status is listed as *`undef`* , meaning undefined. This is different from '-', meaning no value.

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.

| [-health]

Use this parameter to display detailed health information for the specified network ports.

| [-instance] }

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

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

Selects the network ports that match this parameter value. Use this parameter with the `-port` parameter to select a port.

[-port {<netport>|<ifgrp>}] - Port

Selects the network ports that match this parameter value. If you do not use this parameter, the command displays information about all network ports.

[-link {off|up|down}] - Link

Selects the network ports that match this parameter value.

[-mtu <integer>] - MTU

Selects the network ports that match this parameter value.

[-autonegotiate-admin {true|false}] - Auto-Negotiation Administrative

Selects the network ports that match this parameter value.

[-autonegotiate-oper {true|false}] - Auto-Negotiation Operational

Selects the network ports that match this parameter value.

[-duplex-admin {auto|half|full}] - Duplex Mode Administrative

Selects the network ports that match this parameter value.

[-duplex-oper {auto|half|full}] - Duplex Mode Operational

Selects the network ports that match this parameter value.

[-speed-admin {auto|10|100}] - Speed Administrative

Selects the network ports that match this parameter value.

[-speed-oper {auto|10|100}] - Speed Operational

Selects the network ports that match this parameter value.

[-flowcontrol-admin {none|receive|send|full|pfc}] - Flow Control Administrative

Selects the network ports that match this parameter value.

[-flowcontrol-oper {none|receive|send|full|pfc}] - Flow Control Operational

Selects the network ports that match this parameter value.

[-mac <MAC Address>] - MAC Address

Selects the network ports that match this parameter value.

[-up-admin {true|false}] - Up Administrative (privilege: advanced)

Selects the network ports that match this parameter value.

[-type {physical|if-group|vlan|vip|pvlan}] - Port Type

Selects the network ports that match this parameter value.

[-ifgrp-node <nodename>] - Interface Group Parent Node

Selects the network ports that match this parameter value.

[-ifgrp-port {<netport>|<ifgrp>}] - Interface Group Parent Port

Selects the network ports that match this parameter value.

[-ifgrp-distr-func {mac|ip|sequential|port}] - Distribution Function

Selects the network ports that match this parameter value.

[-ifgrp-mode {multimode|multimode_lacp|singlemode}] - Create Policy

Selects the network ports that match this parameter value.

[-vlan-node <nodename>] - Parent VLAN Node

Selects the network ports that match this parameter value.

[-vlan-port {<netport>|<ifgrp>}] - Parent VLAN Port

Selects the network ports that match this parameter value.

[-vlan-tag <integer>] - VLAN Tag

Selects the network ports that match this parameter value.

[-remote-device-id <text>] - Remote Device ID

Selects the network ports that match this parameter value.

[-ipspace <IPspace>] - IPspace Name

Use this parameter to display information only about the ports that match the IPspace you specify.

[-broadcast-domain <Broadcast Domain>] - Broadcast Domain

Use this parameter to display information only about the ports that match the broadcast-domain you specify.

[-mtu-admin <integer>] - MTU Administrative

Selects the network ports that match this parameter value.

[-health-status {healthy|degraded}] - Port Health Status

Use this parameter to display information only about the ports that match the health-status you specify.

[-ignore-health-status {true|false}] - Ignore Port Health Status

Use this parameter to display information only about the ports that match the ignore-health-status you specify.

[-health-degraded-reasons {l2-reachability|link-flapping|crc-errors|vswitch-link}] - Port Health Degraded Reasons

Use this parameter to display information only about the ports that match the degraded-reason you specify.

[-vm-network-name <text>] - Virtual Machine Network Name

Use this parameter to display information only about the ports that match the network name you specify. Google Cloud Platform only.

[-rdma-protocols <roce>,…] - Supported RDMA Protocols

Use this parameter to display information only about the ports that support the specified RDMA protocols.

[-pfc-queues-admin <integer>,…] - List of PFC Queues Administrative

Selects the network ports that match this parameter value.

[-pfc-queues-oper <integer>,…] - List of PFC Queues Operational

Selects the network ports that match this parameter value.

Examples

The following example displays information about all network ports.

```
cluster1::> network port show

Node: node1

Ignore
                                                     Speed(Mbps) Health
Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
--------- ------------ ---------------- ---- ---- ----------- --------
------
e0a       Cluster      Cluster          up   9000  auto/1000  healthy
false
e0b       Cluster      Cluster          up   9000  auto/1000  healthy
false
e0c       Default      Default          up   1500  auto/1000  degraded
false
e0d       Default      Default          up   1500  auto/1000  degraded
true
Node: node2

Ignore
                                                     Speed(Mbps) Health
Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
--------- ------------ ---------------- ---- ---- ----------- --------
------
e0a       Cluster      Cluster          up   9000  auto/1000  healthy
false
e0b       Cluster      Cluster          up   9000  auto/1000  healthy
false
e0c       Default      Default          up   1500  auto/1000  healthy
false
e0d       Default      Default          up   1500  auto/1000  healthy
false

8 entries were displayed.
```

The following example displays health information about all network ports.

```
cluster1::> network port show -health
                                   Ignore
                          Health   Health
Node       Port      Link Status   Status Degraded Reasons
--------   --------- ---- -------- ------ -----------------------
node1
           e0a       up   healthy  false  -
           e0b       up   healthy  false  -
           e0c       up   degraded false  l2_reachability,
                                          link_flapping
           e0d       up   degraded false  l2_reachability

node2
           e0a       up   healthy  false  -
           e0b       up   healthy  false  -
           e0c       up   healthy  false  -
           e0d       up   degraded false  -

8 entries were displayed.
```

network port broadcast-domain add-ports

Add ports to a layer 2 broadcast domain

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

Description

Add ports to a broadcast domain.

> The IPSpace of the ports added will be updated to the IPSpace of the broadcast-domain. The ports will be added to the failover-group of the broadcast-domain. The MTU of the ports will be updated to the MTU of the broadcast-domain.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace of the broadcast domain.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain

The broadcast domain for this port assignment.

`-ports [node>:<port],…` - List of ports

The ports to be added to this broadcast domain.

Examples

The following example adds the port "e0d" on node "cluster1-1" and port "e0d" on node "cluster1-2" to broadcast domain "mgmt" in IPspace "Default".

```
cluster1::network port broadcast-domain> add-ports -ipspace Default
-broadcast-domain mgmt -ports cluster1-1:e0d, cluster1-2:e0d
```

network port broadcast-domain create

Create a new layer 2 broadcast domain

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

Description

Create a new broadcast domain.

> The IPSpace of the ports added will be updated to the IPSpace of the broadcast-domain. A failover-group will be generated containing the ports of the broadcast-domain. The MTU of all of the ports in the broadcast-domain will be updated to the MTU specified for the broadcast-domain.

Parameters

`[-ipspace <IPspace>]` - IPspace Name

The IPspace to which the new broadcast domain belongs.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain

The name of the broadcast domain to be created. The name of the broadcast domain needs to be unique within the IPspace.

`-mtu <integer>` - Configured MTU

MTU of the broadcast domain.

`[-ports [node>:<port],…]` - Ports

The network ports to be added to the broadcast domain. Ports need to be added to the broadcast domain before interfaces can be hosted on the port. By default, no port will be added to the broadcast domain.

Examples

The following example creates broadcast domain "mgmt" in IPspace "Default" with an MTU of 1500 and network ports e0c from node "gx1" and node "gx2".

```
cluster1::> network port broadcast-domain create -ipspace Default
-broadcast-domain mgmt -mtu 1500 -ports gx1:e0c,gx2:e0c
```

network port broadcast-domain delete

Delete a layer 2 broadcast domain

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

Description

Delete a broadcast domain that contains no ports.

Parameters

`-ipspace <IPspace>` - IPspace Name
The IPspace to which the broadcast domain belongs

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain
The name of the broadcast domain to be deleted.

Examples

The following example deletes the broadcast domain "mgmt" in IPspace "Default".

```
cluster1::network port broadcast-domain> delete -ipspace Default
-broadcast-domain mgmt
```

network port broadcast-domain merge

Merges two layer 2 broadcast domains

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

Description

Merges a broadcast domain into an existing broadcast domain.

Parameters

`-ipspace <IPspace>` - IPspace Name
The IPspace of the broadcast domain.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain
The merging broadcast domain.

`-into-broadcast-domain <Broadcast Domain>` - Merge with This Layer 2 Broadcast Domain
The target broadcast domain for the merge operation.

Examples

The following example merges broadcast domain "bd-mgmt" in IPspace "Default" to broadcast domain "bd-data".

```
cluster1::network port broadcast-domain> merge -ipspace Default -broadcast
-domain bd-mgmt -into-broadcast-domain bd-data
```

network port broadcast-domain modify

Modify a layer 2 broadcast domain

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

Description

Modify a broadcast domain.

Parameters

`-ipspace <IPspace>` - IPspace Name
The IPspace to which the broadcast domain belongs.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain
The name of the broadcast domain.

`[-mtu <integer>]` - Configured MTU
MTU of the broadcast domain. All ports that belong to the broadcast domain are modified.

Examples

The following example modifies the mtu attribute of broadcast domain "mgmt" in IPspace "Default" to 1500

```
cluster1::network port broadcast-domain*> modify -ipspace Default
-broadcast-domain mgmt -mtu 1500
```

network port broadcast-domain move

Move a layer 2 broadcast domain to another IPspace

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

Description

Move a broadcast domain to another IPspace.

Parameters

`-ipspace <IPspace>` - IPspace Name
The IPspace to which the broadcast domain belongs.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain

The name of the broadcast domain.

`-to-ipspace <IPspace>` - Name of the destination IPspace

The name of the destination IPspace.

Examples

The following example moves the broadcast domain named "mgmt" from IPspace "Default" to IPspace "Default-1".

```
cluster1::network port broadcast-domain> move -ipspace Default -broadcast
-domain mgmt -to-ipspace Default-1
```

network port broadcast-domain remove-ports

Remove ports from a layer 2 broadcast domain

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

Description

Remove port assignments from a broadcast domain.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace of the broadcast domain.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain

The broadcast domain of the ports.

`-ports [node>:<port],…` - List of ports

The ports to removed from the broadcast-domain.

Examples

The following example removes port "e0d" on node "cluster1-1" and port "e0d" on node "cluster1-2" from broadcast domain "mgmt" in IPspace "Default".

```
cluster1::network port broadcast-domain> remove-ports -ipspace Default
-broadcast-domain mgmt -ports cluster1-1:e0d, cluster1-2:e0d
```

network port broadcast-domain rename

Rename a layer 2 broadcast domain

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

Description

Rename a broadcast domain.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace to which the broadcast domain belongs.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain

The name of the broadcast domain.

`-new-name <text>` - New Name

The new name of the broadcast domain.

Examples

The following example renames the broadcast domain named "mgmt" to "mgmt2" in IPspace "Default".

```
cluster1::network port broadcast-domain> rename -ipspace Default
-broadcast-domain mgmt -new-name mgmt2
```

network port broadcast-domain show

Display layer 2 broadcast domain information

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

Description

Display broadcast domain information.

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.

`[-ipspace <IPspace>]` - IPspace Name

Selects the broadcast domains that match the IPspace name.

`[-broadcast-domain <Broadcast Domain>]` - Layer 2 Broadcast Domain

Selects the broadcast domains that match the broadcast domain name.

`[-mtu <integer>]` - Configured MTU

Selects the broadcast domains that match the MTU value. This field is the MTU that was configured by the

user, which might be different from the operational MTU.

[-ports [node>:<port],…] - Ports

Selects the broadcast domains that contain the network ports. For example, node1:e0a will display broadcast domains that contain node1:e0a network port.

[-port-update-status {complete|in-progress|error|overridden-while-offline}] - Port Update Status

Selects the broadcast domains that contain the network port status. For example, specifying "error" will display broadcast domains that contain "Error" network port status.

[-port-update-status-details <text>,…] - Status Detail Description

Selects the broadcast domains that contain the network port status detail text.

[-port-update-status-combined {complete|in-progress|error|overridden-while-offline}] - Combined Port Update Status

Selects the broadcast domains that contain the combined network port status. For example, specifying "error" will display broadcast domains that contain a combined network port status of "Error".

[-failover-groups <failover-group>,…] - Failover Groups

Selects the broadcast domains that contain the failover groups.

[-subnet-names <subnet name>,…] - Subnet Names

Selects the broadcast domains that contain the subnet name or names.

[-is-vip {true|false}] - Is VIP Broadcast Domain

Selects the broadcast domains that match a specific "is-vip" flag. Specifying "true" matches only broadcast domains associated with the Virtual IP feature; "false" matches only broadcast domains that do not.

Examples

The following example displays general information about broadcast domains.

```
cluster1::> network port broadcast-domain show
IPspace Broadcast                                          Update
Name    Domain Name    MTU  Port List                      Status Details
------- ----------- ------  ----------------------------- --------------
Cluster Cluster       9000
                            node1:e0a                     complete
                            node1:e0b                     complete
Default Default       1500
                            node1:e0c                     complete
                            node1:e0d                     complete
2 entries were displayed.
```

network port broadcast-domain split

Splits a layer 2 broadcast domain into two layer 2 broadcast domains.

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

Description

Splits ports from a broadcast domain into a new broadcast domain.

The following restrictions apply to this command:

- If the ports are in a failover group, all ports in the failover group must be provided. Use network interface failover-groups show to see which ports are in failover groups.
- If the ports have LIFs associated with them, the LIFs cannot be part of a subnet's ranges and the LIF's `curr-port` and `home-port` must both be provided. Use network interface show`-fields`*subnet-name* , *home-node* , *home-port* , *curr-node* , *curr-port* to see which ports have LIFs associated with them and whether the LIFs are part of a subnet's ranges. Use `network subnet remove-ranges` with the LIF's IP address and `-force-update-lif-associations` set to *true* to remove the LIF's association with a subnet.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace of the broadcast domain.

`-broadcast-domain <Broadcast Domain>` - Layer 2 Broadcast Domain

The broadcast domain to split.

`-new-broadcast-domain <Broadcast Domain>` - New Layer 2 Broadcast Domain Name

The new broadcast domain.

`-ports [node>:<port],…` - List of Ports

The ports to be split from this broadcast domain.

Examples

The following example splits port "e0d" on node "cluster1-1" and port "e0d" on node "cluster1-2" from broadcast domain "bd-mgmt" in IPspace "Default" to broadcast domain "bd-data".

```
cluster1::> network port broadcast-domain split -ipspace Default
-broadcast-domain bd-mgmt -new-broadcast-domain bd-data -ports cluster1-
1:e0d, cluster1-2:e0d
```

Related Links

- network interface failover-groups show
- network interface show
- network subnet remove-ranges

network port ifgrp add-port

Add a port to an interface group

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

Description

The `network port ifgrp add-port` command adds a network port to a port interface group. The port interface group must already exist. You can create a port interface group by using the network port ifgrp create command.

The following restrictions apply to port interface groups:

- A port that is already a member of a port interface group cannot be added to another port interface group.
- Cluster ports and management ports cannot be in a port interface group.
- A port to which a logical interface is already bound cannot be added to a port interface group.
- A port that already has an assigned failover role cannot be added to a port interface group.
- A VLAN port cannot be added to a port interface group.
- A port which attaches to a VLAN cannot be added to a port interface group.
- An interface group port cannot be added to a port interface group.
- A port that is assigned to a broadcast domain cannot be added to a port interface group.
- All ports in a port interface group must be physically located on the same node.

Parameters

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

The node on which the port interface group is located.

`-ifgrp <ifgrp name>` - Interface Group Name

The port interface group to which a port is to be added.

`-port <netport>` - Specifies the name of port.

The network port that is to be added to the port interface group.

`[-skip-broadcast-domain-placement <true>]` - Skip Placement Into Broadcast Domain (privilege: advanced)

When specified along with the first port added to the ifgrp, the ifgrp will not be placed into any broadcast domain.

Examples

The following example adds port e0c to port interface group a1a on a node named node1:

```
cluster1::> network port ifgrp add-port -node node1 -ifgrp a1a -port e0c
```

Related Links

- network port ifgrp create

network port ifgrp create

Create a port interface group

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

Description

The `network port ifgrp create` command creates a port interface group. See the documentation for the network port ifgrp add-port command for a list of restrictions on creating port interface groups.

Parameters

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

The node on which the port interface group will be created.

`-ifgrp <ifgrp name>` - Interface Group Name

The name of the port interface group that will be created. Port interface groups must be named using the syntax "a<number><letter>", where <number> is an integer in the range [0-999] without leading zeros and <letter> is a lowercase letter. For example, "a0a", "a0b", "a1c", and "a2a" are all valid port interface group names.

`-distr-func {mac|ip|sequential|port}` - Distribution Function

The distribution function of the port interface group that will be created. Valid values are:

- mac - Network traffic is distributed based on MAC addresses
- ip - Network traffic is distributed based on IP addresses
- sequential - Network traffic is distributed in round-robin fashion from the list of configured, available ports
- port - Network traffic is distributed based on the transport layer (TCP/UDP) ports

`-mode {multimode|multimode_lacp|singlemode}` - Create Policy

The create policy for the interface group that will be created. Valid values are:

- multimode - Bundle multiple member ports of the interface group to act as a single trunked port
- multimode_lacp - Bundle multiple member ports of the interface group using Link Aggregation Control Protocol
- singlemode - Provide port redundancy using member ports of the interface group for failover

Examples

The following example creates a port interface group named a0a on node node0 with a distribution function of ip:

```
cluster1::> network port ifgrp create -node node0 -ifgrp a0a -distr-func
ip -mode multimode
```

Related Links

- network port ifgrp add-port

network port ifgrp delete

Destroy a port interface group

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

Description

The `network port ifgrp delete` command destroys a port interface group.

> When you delete an interface group port, it is automatically removed from failover rules and groups to which it belongs.

Parameters

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

The node on which the port interface group is located.

`-ifgrp <ifgrp name>` - Interface Group Name

The port interface group that will be deleted.

Examples

The following example deletes port interface group a0b from a node named node0.

```
cluster1::> network port ifgrp delete -node node0 -ifgrp a0b
```

network port ifgrp remove-port

Remove a port from an interface group

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

Description

The `network port ifgrp remove-port` command removes a network port from a port interface group.

Parameters

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

The node on which the port interface group is located.

`-ifgrp <ifgrp name>` - Interface Group Name

The port interface group from which a port will be removed.

`-port <netport>` - Specifies the name of port.

The network port that will be removed from the port interface group.

Examples

The following example removes port e0d from port interface group a1a on a node named node1:

```
cluster1::> network port ifgrp remove-port -node node1 -ifgrp a1a -port
e0d
```

network port ifgrp show

Display port interface groups

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

Description

The `network port ifgrp show` command displays information about port interface groups. By default, it displays information about all port interface groups on all nodes in the cluster.

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>|local}] - Node

Selects the port interface groups that match this parameter value. Use this parameter with the `-ifgrp` parameter to select information about a specific port interface group.

[-ifgrp <ifgrp name>] - Interface Group Name

Selects the port interface groups that match this parameter value. Use this parameter with the `-node` parameter, to select information about a specific port interface group.

[-distr-func {mac|ip|sequential|port}] - Distribution Function

Selects the port interface groups that match this parameter value.

[-mode {multimode|multimode_lacp|singlemode}] - Create Policy

Selects the port interface groups that match this parameter value.

[-mac <MAC Address>] - MAC Address

Selects the port interface groups that match this parameter value.

[-activeports {full|partial|none}] - Port Participation

Selects the port interface groups that match this parameter value. The value "partial" indicates that some but not all of the port interface group's ports are active. the value "full" indicates that all of the port interface group's ports are active.

[-ports {<netport>|<ifgrp>}] - Network Ports

Selects the port interface groups that match this parameter value.

[-up-ports {<netport>|<ifgrp>}] - Up Ports

Selects the port interface groups that match this parameter value. Displays only the ports that are up.

[-down-ports {<netport>|<ifgrp>}] - Down Ports

Selects the port interface groups that match this parameter value. Displays only the ports that are down.

Examples

The following example displays information about all port interface groups.

```
cluster1::> network port ifgrp show
        Port     Distribution                Active
Node    ifgrp    Function     MAC Address    Ports   Ports
------  -------  ------------ -------------- ------- ---------------
node0
        a0a      ip           b8:f8:7a:20:00 partial e0c
node1
        a1a      ip           07:26:60:02:00 full    e0d
```

network port reachability repair

Repair reachability

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

Description

Repair network port configuration to match the detected reachable broadcast domain. If the reachability scan has completed with a reachability-status of misconfigured-reachability, meaning the network port does not have reachability to its configured broadcast domain, but has reachability to another broadcast domain, then repairing the network port's reachability will assign the port to the detected broadcast domain. Similarly, if the reachability scan has completed with a reachability status of no-reachability, then repairing the network port's reachability will assign the port to an empty broadcast domain. LIFs configured on the port will be adjusted to be configured on another port in their current broadcast domain if possible. Vlans on the specified port that do not have reachability to their configured broadcast domain will be removed. If the port was part of an ifgrp, the port will be removed from the ifgrp. If the port is not configured on a broadcast domain and has no reachability to any existing broadcast domains, it will be configured in a newly created broadcast domain.

Parameters

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

Selects the node on which the port resides.

`-port <netport>` - Port

Selects the port on which to repair configuration.

Examples

The following example applies the scanned broadcast domain reachability information to the specified port's configuration.

```
cluster1::> network port reachability repair -node node1 -port e0c
```

network port reachability scan

Perform a reachability scan

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

Description

Scan network port configuration to detect reachable broadcast domains.

Parameters

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

Selects the node on which the port resides.

`-port <netport>` - Port

Selects the port on which to scan broadcast domain reachability.

Examples

The following example applies the scanned broadcast domain reachability information to the specified port's configuration.

```
cluster1::> network port reachability scan -node node1 -port e0c
```

network port reachability show

Display Reachability Status

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

Description

Display broadcast domain reachability information for the specified port. If the port is discovered, via a layer 2 reachability scan, to have reachability to broadcast domains other than the one on which it is expected, the

command will list the reachable broadcast domains and an appropriate reachability status message. If the reachable broadcast domain matches the expected one, the reachability status is displayed as Ok.

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.

| [-detail]

When true, additional details regarding which broadcast domains have been found to be reachable from the specified network port are displayed.

| [-instance] }

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

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

Selects the node on which the port resides.

[-port <netport>] - Port

Selects the port for which to display broadcast domain reachability information.

[-expected-broadcast-domain [IPspace>:<Broadcast Domain]] - Expected Broadcast Domain

The IPspace and broadcast domain currently assigned to the network port. If the specified port is an ifgrp member, the expected broadcast domain is the IPspace and broadcast domain currently assigned to the parent ifgrp.

[-reachable-broadcast-domains [IPspace>:<Broadcast Domain],…] - Reachable Broadcast Domains

The name of the IPspaces and broadcast domains that have reachability to the specified port, as discovered by a layer 2 reachability scan.

[-reachability-status {unknown|ok|no-reachability|misconfigured-reachability|multi-domain-reachability|internal-error}] - Reachability Status

The status of the broadcast domain reachability for the specified port. "Ok" if the expected broadcast domain matches the reachable broadcast domains, i.e., the port can reach other ports in the expected broadcast domain, but no ports configured in other broadcast domains. "No-reachability" if the port cannot reach any ports in the expected broadcast domain, and also cannot reach any ports in any other broadcast domains. "Misconfigured-reachability" if the port cannot reach any ports in the expected broadcast domain, but can reach ports in one other broadcast domain. "Multi-domain-reachability" if the port can reach other ports configured in multiple broadcast domains. "Unknown" if the port has not been link-up long enough for reachability to be determined.

[-unreachable-ports [node>:<port],…] - Unreachable Ports

The list of network ports that are expected in the same broadcast domain as the specified port but cannot be reached, either because those ports are down or because there is no network connectivity to those ports.

[-unexpected-ports [node>:<port],…] - Unexpected Ports

The list of network ports that are not expected in the same broadcast domain yet have network connectivity to the specified port.

Examples

The following example displays the broadcast domain reachability for the specified port.

```
cluster1::> network port reachability show -node node1 -port e0d
network port reachability show)
ode         Port     Expected Reachability        Reachability Status
----------- -------- ----------------------------
--------------------------
ode1        e0d      Default:Default              ok
```

The following example displays the detailed reachability for the 'e0d' port when it has misconfigured-reachability, i.e., it cannot reach the other 'e0d' port in the expected broadcast domain 'Default:Default', but can reach the 'e0c' ports configured in the 'Default:Default-3' broadcast domain.

```
cluster1::> network port reachability show -node node1 -port e0d -detail
network port reachability show)
ode         Port     Expected Reachability        Reachability Status
----------- -------- ----------------------------
--------------------------
ode1        e0d      Default:Default              misconfigured-
reachability
Actual Reachability: Default:Default-3
Unreachable Ports: node2:e0d
 Unexpected Ports: node2:e0c, node1:e0c
```

The following example displays the detailed reachability for the 'e0d' port when it has multi-domain-reachability, i.e., it can reach the other 'e0d' port in the expected broadcast domain 'Default:Default', but can also reach the 'e0c' ports configured in the 'Default:Default-3' broadcast domain.

```
cluster1::> network port reachability show -node node1 -port e0d -detail
network port reachability show)
ode         Port     Expected Reachability        Reachability Status
----------- -------- ----------------------------
--------------------------
ode1        e0d      Default:Default              multi-domain-
reachability
Actual Reachability: Default:Default,
                     Default:Default-3
Unreachable Ports: -
 Unexpected Ports: node2:e0c, node1:e0c
```

The following example displays the detailed reachability for the 'e0d' port when it has no-reachability, i.e., it cannot reach the other 'e0d' port in the expected broadcast domain 'Default:Default', and also cannot reach any other ports configured in broadcast domains.

```
cluster1::> network port reachability show -node node1 -port e0d -detail
network port reachability show)
ode         Port     Expected Reachability        Reachability Status
----------- -------- ----------------------------
--------------------------
ode1        e0d      Default:Default              no-reachability
Actual Reachability: -
Unreachable Ports: node2:e0d
 Unexpected Ports: -
```

network port vip create

Create a VIP port

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

Description

The `network port vip create` command creates a VIP port in the specified IPspace on the specified node. Only one VIP port can be created per IPspace on the given node.

Parameters

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

The node where the VIP port should be created.

`-port <netport>` - Network Port

The name of the VIP port to be created in the format v<slot-number><port-letter>

`-ipspace <IPspace>` - IPspace Name

The IPspace where the VIP port should be created. The default value for this parameter is "Default", which identifies the default IPspace.

Examples

This example shows how to create a VIP port named v0a in ipspace ips on node1.

```
cluster1::> network port vip create -node node1 -port v0a -ipspace ips
```

network port vip delete

Delete a VIP port

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

Description

The `network port vip delete` command deletes a VIP port.

Parameters

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

The node associated with the VIP port to be deleted.

`-port <netport>` - Network Port

The name of the VIP port to be deleted.

Examples

This example shows how to delete VIP Port v0a on node1.

```
cluster1::> network port vip delete -node node1 -port v0a
```

network port vip show

Display VIP ports

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

Description

The `network port vip show` command displays information about VIP ports.

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>|local}] - Node

This parameter selects the VIP ports that match the specified node.

[-port <netport>] - Network Port

This parameter selects the VIP ports that match the specified port.

[-ipspace <IPspace>] - IPspace Name

This parameter selects the VIP ports that match the specified IPspace.

Examples

The example below shows VIP port v0a is created in IPspace ips on node1.

```
cluster1::> network port vip show
    Node   VIP Port IPspace
    ------ -------- -------------
    node1  v0a      ips
```

network port vlan create

Create a virtual LAN (VLAN)

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

Description

The `network port vlan create` command attaches a VLAN to a network port on a specified node.

Parameters

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

The node to which the VLAN is to be attached.

> You cannot attach a VLAN to a cluster port.

{ `-vlan-name {<netport>|<ifgrp>}` - VLAN Name

The name of the VLAN that is to be attached. This name should be a combination of the name of the port or interface group and the VLAN ID, with a hyphen between, such as "e1c-80".

| `-port {<netport>|<ifgrp>}` - Associated Network Port

The network port to which the VLAN is to be attached.

`-vlan-id <integer>` - Network Switch VLAN Identifier }

The ID tag of the created VLAN.

`[-skip-broadcast-domain-placement <true>]` - Skip Placement Into Broadcast Domain (privilege: advanced)

When specified, the VLAN will not be placed into any broadcast domain.

Examples

This example shows how to create VLAN e1c-80 attached to network port e1c on node1.

```
cluster1::> network port vlan create -node node1 -vlan-name e1c-80
```

network port vlan delete

Delete a virtual LAN (VLAN)

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

Description

The `network port vlan delete` command deletes a VLAN from a network port.

> When you delete a VLAN port, it is automatically removed from all failover rules and groups that use it.

Parameters

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

The node from which the VLAN is to be deleted.

{ `-vlan-name {<netport>|<ifgrp>}` - VLAN Name

The name of the VLAN that is to be deleted

| `-port {<netport>|<ifgrp>}` - Associated Network Port

The network port to which the VLAN is to be attached.

`-vlan-id <integer>` - Network Switch VLAN Identifier }

The ID tag of the deleted VLAN.

Examples

This example shows how to delete VLAN e1c-80 from network port e1c on node1.

```
cluster1::> network port vlan delete -node node1 -vlan-name e1c-80
```

network port vlan show

Display virtual LANs (VLANs)

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

Description

The `network port vlan show` command displays information about network ports that are attached to VLANs. The command output indicates any inactive links and lists the reason for the inactive status.

If the operational duplex mode and speed cannot be determined (for instance, if the link is down), they are listed as `undef` , meaning undefined.

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>|local}] - Node

Selects the VLAN network ports that match this parameter value.

{ [-vlan-name {<netport>|<ifgrp>}] - VLAN Name

Selects the VLAN network ports that match this parameter value.

| [-port {<netport>|<ifgrp>}] - Associated Network Port

Selects the VLAN network ports that match this parameter value. If neither this parameter nor `-name` are used, the command displays information about all network ports.

[-vlan-id <integer>] - Network Switch VLAN Identifier }

Selects the VLAN network ports that match this parameter value.

[-mac <MAC Address>] - MAC address

Selects the VLAN network ports that match this parameter value.

Examples

The example below shows VLAN e1b-70 attached to port e1b on node1.

```
cluster1::> network port vlan show
                 Network Network
Node   VLAN Name Port    VLAN ID  MAC Address
------ --------- ------- -------  ------------------
node1  e1b-70    e1b     70       00:15:17:76:7b:69
```

network qos-marking commands

network qos-marking modify

Modify the QoS marking values

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

Description

The `network qos-marking modify` command modifies the QoS marking values for different protocols, for each IPspace.

Parameters

-ipspace <IPspace> - IPspace name

Use this parameter to specify the IPspace for which the QoS marking entry is to be modified.

-protocol <text> - Protocol

Use this parameter to specify the protocol for which the QoS marking entry is to be modified. The possible values are NFS, CIFS, iSCSI, NVMe-TCP, SnapMirror, SnapMirror-Sync, NDMP, FTP, HTTP-admin, HTTP-filesrv, SSH, Telnet, and SNMP.

[-dscp <integer>] - DSCP Marking Value

Use this parameter to specify the DSCP value. The possible values are 0 to 63.

[-is-enabled {true|false}] - Is QoS Marking Enabled

Use this parameter to enable or disable the QoS marking for the specified protocol and IPspace.

Examples

The following example modifies the QoS marking entry for the NFS protocol in the Default IPspace:

```
cluster1::> network qos-marking modify -ipspace Default -protocol NFS
-dscp 10 -is-enabled true
```

network qos-marking show

Display the QoS marking values

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

Description

The `network qos-marking show` command displays the QoS marking values for different protocols, for each IPspace.

Parameters

{ [-fields <fieldname>,…]

Use this parameter to display only certain fields of the QoS marking table.

| [-instance] }

Use this parameter to display all the fields of the QoS marking table.

[-ipspace <IPspace>] - IPspace name

Use this parameter to display the QoS marking entries for the specified IPspace.

[-protocol <text>] - Protocol

Use this parameter to display the QoS marking entries for the specified protocol. The possible values are NFS, CIFS, iSCSI, NVMe-TCP, SnapMirror, SnapMirror-Sync, NDMP, FTP, HTTP-admin, HTTP-filesrv, SSH, Telnet, and SNMP.

[-dscp <integer>] - DSCP Marking Value

Use this parameter to display the QoS marking entries matching the specified DSCP value. The possible values are 0 to 63.

[-is-enabled {true|false}] - Is QoS Marking Enabled

Use this parameter to display the QoS marking entries matching the specified flag.

Examples

The following example displays the QoS marking entries for the Default IPspace.

```
cluster1::> network qos-marking show -ipspace Default
IPspace             Protocol           DSCP  Enabled?
------------------- ----------------- -----  --------
Default
                    CIFS                 10  false
                    FTP                  48  false
                    HTTP-admin           48  false
                    HTTP-filesrv         10  false
                    NDMP                 10  false
                    NFS                  10  true
                    SNMP                 48  false
                    SSH                  48  false
                    SnapMirror           10  false
                    SnapMirror-Sync      10  false
                    Telnet               48  false
                    iSCSI                10  false
                    NVMe-TCP             10  false
12 entries were displayed.
```

network route commands

network route create

Create a static route

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

Description

The `network route create` command creates a static route within a Vserver.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the Vserver on which the route is to be created.

`-destination <IP Address/Mask>` - Destination/Mask

Use this parameter to specify the IP address and subnet mask of the route's destination. The format for this value is: address, slash ("/"), mask. `0.0.0.0/0` is a valid destination value to create default IPv4 route. And `::/0` is a valid destination value to create default IPv6 route

`-gateway <IP Address>` - Gateway

Use this parameter to specify the IP address of the gateway server leading to the route's destination.

[-metric <integer>] - Metric

Use this parameter to specify the metric of the route.

Examples

The following example creates default routes within Vserver vs0 for IPv4 and IPv6.

```
cluster1::> network route create -vserver vs0 -destination 0.0.0.0/0
-gateway 10.61.208.1
cluster1::> network route create -vserver vs0 -destination ::/0 -gateway
3ffe:1::1
```

network route delete

Delete a static route

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

Description

The `network route delete` command deletes a static route from a Vserver.

Parameters

-vserver <vserver> - Vserver Name

Use this parameter to specify the Vserver on which the route is to be deleted.

-destination <IP Address/Mask> - Destination/Mask

Use this parameter to specify the IP address and subnet mask of the route's destination. The format for this value is: address, slash ("/"), mask. For example, `0.0.0.0/0` is a correctly formatted value for the -destination parameter.

-gateway <IP Address> - Gateway

Use this parameter to specify the gateway on which the route is to be deleted.

Examples

The following example deletes a route within Vserver vs0 for destination 0.0.0.0/0.

```
cluster1::network route delete -vserver vs0 -destination 0.0.0.0/0
```

network route show-lifs

Show the Logical Interfaces for each route entry

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

Description

The `network route show-lifs` command displays the association of static routes and Logical Interfaces (LIFs) within one or more Vservers. You can view routes within specified Vservers, routes with specified destinations, and routes with specified gateways.

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.

[-vserver <vserver>] - Vserver Name

Use this parameter to display only routes that have the specified Vserver as their origin.

[-destination <IP Address/Mask>] - Destination/Mask

Use this parameter to display only routes that have the specified IP address and subnet mask as their destination. The format for this value is: address, slash ("/"), mask. For example, `0.0.0.0/0` is a valid value for the -destination parameter.

[-gateway <IP Address>] - Gateway

Use this parameter to display only routes that have the specified IP address as their gateway.

[-lifs <lif-name>,…] - Logical Interfaces

Use this parameter to display only the routes that are associated with the specified Logical Interfaces (LIFs).

[-address-family {ipv4|ipv6|ipv6z}] - Address Family

Use this parameter to display only the routes that belong to specified address family.

network route show

Display static routes

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

Description

The `network route show` command displays a group of static routes within one or more Vservers. You can view routes within specified Vservers, routes with specified destinations, and routes with specified gateways.

Parameters

{ [-fields <fieldname>,…]

Use this parameter to display only certain fields of the routing tables.

| [-instance] }

Use this parameter to display all fields of the routing tables.

[-vserver <vserver>] - Vserver Name

Use this parameter to display only routes that have the specified Vserver as their origin.

[-destination <IP Address/Mask>] - Destination/Mask

Use this parameter to display only routes that have the specified IP address and subnet mask as their destination. The format for this value is: address, slash ("/"), mask. The example below has `0.0.0.0/0` as a valid value for the -destination parameter.

[-gateway <IP Address>] - Gateway

Use this parameter to display only routes that have the specified IP address as their gateway.

[-metric <integer>] - Metric

Use this parameter to display only routes that have the specified metric.

[-ipspace <IPspace>] - IPspace Name

Use this parameter to optionally specify the IPspace associated with the Vserver. This parameter can be used in conjunction with the Vserver parameter in order to configure the same route across multiple Vservers within an IPspace.

[-address-family {ipv4|ipv6|ipv6z}] - Address family of the route

Use this parameter to display only the routes that have the specified address-family.

Examples

The following example displays information about all routing groups.

```
cluster1::> network route show
  (network route show)
Server              Destination     Gateway         Metric
------------------- --------------- --------------- ------
node1
                    0.0.0.0/0       10.61.208.1     20
node2
                    0.0.0.0/0       10.61.208.1     20
vs0
                    0.0.0.0/0       10.61.208.1     20
3 entries were displayed.
```

network route active-entry show

Display active routes

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

Description

The `network route active-entry show` command displays installed routes on one or more nodes. You can view routes within specified nodes, within specified Vservers, routes in specified subnet groups, and routes with specified destinations.

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.

| `[-verbose ]` (privilege: advanced)

Use this parameter to display the reference count, use, interface, and Path MTU fields.

| `[-instance ]` (privilege: advanced) }

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

`[-vserver <vserver>]` - Vserver Name (privilege: advanced)

Displays the routes that have the specified Vserver as their origin.

`[-node {<nodename>|local}]` - Node (privilege: advanced)

Displays the routes from the specified node.

`[-address-type {ipv4|ipv6|ipv6z}]` - Address Family (privilege: advanced)

Displays the routes that have the specified IP address type.

`[-subnet-group <IP Address/Mask>]` - Subnet Group (privilege: advanced)

Displays the group of routes that belong to the specified subnet. Routes within the specified subnet group are used first before the default set. The "default" subnet group is a system-provided set of default routes.

`[-destination <text>]` - Destination (privilege: advanced)

Displays the routes that have the specified IP address or subnet as their destination. The format for the subnet is: <address>/<mask>. IPv6 address includes the scope value after percentage ("%"). 0.0.0.0/0, 169.254.4.60, ff02::%e0a/32 and fe80::250:56ff:fea6:db7c%e0b are valid values for this parameter.

`[-interface <text>]` - Interface Name (privilege: advanced)

Displays the routes that use the specified interface to transmit packets to the destination. A valid interface has the format of {<netport>|<ifgrp>}, such as "e0a", "e0a-1" and "a0a", or it can be a loopback interface, such as "lo" and "losk".

`[-route-interface-address <text>]` - Route Interface Address (privilege: advanced)

Displays the routes that use the specified IP address on the transmit interface.

`[-gateway <text>]` - Route gateway (privilege: advanced)

Displays the routes that have the specified gateway. The gateway can be an IP address, such as "10.10.2.1" and "fe80::1%lo", MAC address, such as "0:5:73:a0:1:7" or refer to a local link, such as "link#3".

`[-metric <integer>]` - Metric (privilege: advanced)

Displays the routes that have the specified metric.

[-flags <text>] - Flags (privilege: advanced)

Displays the routes that have the specified flags. The type string for "-flags" needs to be one or more of the following {U|G|H|R|D|S|1|2|L|C} in the order shown.

- U - Usable
- G - Gateway
- H - Host
- R - Reject
- D - Dynamic
- S - Static
- 1 - Protocol1
- 2 - Protocol2
- L - Llinfo
- C - Clone

Multiple values can be specified (for example: UHL).

[-reference-count <integer>] - Reference Count (privilege: advanced)

Displays the routes that have the specified reference count in the system.

[-lookup-count <integer>] - Lookup Count (privilege: advanced)

Displays the routes that have the specified use count (the count of lookups for the route).

[-path-mtu <integer>] - Path MTU (privilege: advanced)

Displays the routes that have the specified path maximum transmission unit.

Examples

The following example displays active routes on all nodes in Vserver vs0 with subnet-group 10.10.10.0/24.

```
cluster1::*> network route active-entry show -vserver vs0 -subnet-group
10.10.10.0/24
  (network route active-entry show)

Vserver: vs0
Node: node1
Subnet Group: 10.10.10.0/24
Destination             Gateway              Interface   Metric  Flags
----------------------  -------------------  ---------   ------  -----
default                 10.10.10.1           e0c              0  UGS

Vserver: vs0
Node: node2
Subnet Group: 10.10.10.0/24
Destination             Gateway              Interface   Metric  Flags
----------------------  -------------------  ---------   ------  -----
default                 10.10.10.1           e0c              0  UGS
2 entries were displayed.
```

network subnet commands

network subnet add-ranges

Add new address ranges to a subnet

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

Description

Add new address ranges to a subnet.

> All addresses in a range must be the same address family (IPv4 or IPv6) and must have the same subnet mask. Ranges that overlap or are next to existing ranges will be merged with the existing ranges.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace in which the range resides.

`-subnet-name <subnet name>` - Subnet Name

The name of the subnet.

`-ip-ranges {<ipaddr>|<ipaddr>-<ipaddr>}` - IP Ranges

The list of ranges to add to the subnet.

[-force-update-lif-associations <true>] - Force Update LIF Associations

This command will fail if any service processor interfaces or network interfaces are using the IP addresses in the ranges provided. Using this parameter will associate any manually addressed interfaces with the subnet and will allow the command to succeed.

Examples

The following example allocates addresses for subnet *s1* in IPspace *Default* .

```
cluster1::> network subnet add-ranges -ipspace Default -subnet-name s1
    -ip-ranges "10.98.1.20-10.98.1.30, 10.98.1.35, 10.98.1.40-10.98.1.49"
```

network subnet create

Create a new layer 3 subnet

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

Description

Create a new subnet.

Parameters

[-ipspace <IPspace>] - IPspace Name

The IPspace to which the new subnet belongs.

-subnet-name <subnet name> - Subnet Name

The name of the subnet to be created. The name of the subnet needs to be unique within the IPspace.

-broadcast-domain <Broadcast Domain> - Broadcast Domain

The broadcast domain to which the new subnet belongs.

-subnet <IP Address/Mask> - Layer 3 Subnet

The address and mask of the subnet.

[-gateway <IP Address>] - Gateway

The gateway of the subnet.

[-ip-ranges {<ipaddr>|<ipaddr>-<ipaddr>}] - IP Addresses or IP Address Ranges

The IP ranges associated with this subnet.

[-force-update-lif-associations <true>] - Change the subnet association

This command will fail if any service processor interfaces or network interfaces are using the IP addresses in the ranges provided. Using this parameter will associate any manually addressed interfaces with the subnet and will allow the command to succeed.

Examples

The following examples create subnets named *s1* and *s6* in IPspace *Default* .

```
cluster1::> network subnet create -ipspace Default -broadcast-domain bd1
-subnet-name s1
    -subnet 192.168.1.0/24 -gateway 192.168.1.1 -ip-ranges "192.168.1.1-
192.168.1.100, 192.168.1.112, 192.168.1.145"
```

```
cluster1::> network subnet create -ipspace Default -broadcast-domain bd1
-subnet-name s6
    -subnet 3FFE::/64 -gateway 3FFE::1 -ip-ranges "3FFE::10-3FFE::20"
```

network subnet delete

Delete an existing subnet object

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

Description

Delete a subnet that contains no ports.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace to which the subnet belongs.

`-subnet-name <subnet name>` - Subnet Name

The name of the subnet to be deleted.

`[-force-update-lif-associations <true>]` - Change the subnet association

This command will fail if the subnet has ranges containing any existing service processor interface or network interface IP addresses. Setting this value to true will remove the network interface associations with the subnet and allow the command to succeed. However, it will not affect service processor interfaces.

Examples

The following example deletes subnet *s1* in IPspace *Default* .

```
cluster1::> network subnet delete -ipspace Default -subnet-name s1
```

network subnet modify

Modify a layer 3 subnet

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

Description

Modify a subnet.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace to which the subnet belongs.

`-subnet-name <subnet name>` - Subnet Name

The name of the subnet to modify.

`[-subnet <IP Address/Mask>]` - Layer 3 Subnet

The new address and mask of the subnet.

`[-gateway <IP Address>]` - Gateway

The new gateway address.

`[-ip-ranges {<ipaddr>|<ipaddr>-<ipaddr>}]` - IP Addresses or IP Address Ranges

The new IP ranges for this subnet.

`[-force-update-lif-associations <true>]` - Change the subnet association

This command will fail if any existing service processor interfaces or network interfaces are using IP addresses in the IP ranges being added. It will also fail if any existing service processor interfaces or network interfaces are using IP addresses in the IP ranges being removed. Using this parameter will associate the interfaces with the IP addresses in the ranges being added to the subnet. It will also remove the subnet's association with the interfaces with IP addresses in the IP ranges being removed and will allow the command to succeed.

Examples

The following example modifies the subnet address and gateway.

```
cluster1::> network subnet modify -ipspace Default -subnet-name s1 -subnet
192.168.2.0/24 -gateway 192.168.2.1
```

network subnet remove-ranges

Remove address ranges from a subnet

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

Description

Remove address ranges from a subnet.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace in which the range resides.

`-subnet-name <subnet name>` - Subnet Name

The name of the subnet.

`-ip-ranges {<ipaddr>|<ipaddr>-<ipaddr>}` - IP Ranges

IP ranges to remove.

`[-force-update-lif-associations <true>]` - Force Update LIF Associations

This command will fail if any existing service processor interfaces or network interfaces are using IP addresses in the ranges provided. Using this parameter will remove the subnet's association with those interfaces and allow the command to succeed.

Examples

The following example removes an address range with starting address of *`10.98.1.1`* from subnet *`s1`* in IPspace *`Default`* .

```
cluster1::> network subnet remove-ranges -ipspace Default -subnet-name s1
    -ip-ranges "10.98.1.1-10.98.1.30"
```

network subnet rename

Rename a layer 3 subnet

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

Description

Rename a Subnet.

Parameters

`-ipspace <IPspace>` - IPspace Name

The IPspace to which the subnet belongs.

`-subnet-name <subnet name>` - Subnet Name

The name of the subnet to rename.

`-new-name <text>` - New Name

The new name for the subnet.

Examples

The following example renames subnet *`s1`* to *`s3`* .

```
cluster1::> network subnet rename -ipspace Default -subnet s1 -new-name s3
```

network subnet show

Display subnet information

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

Description

Display subnet information.

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.

[-ipspace <IPspace>] - IPspace Name

Selects the subnets that match the given IPspace name.

[-subnet-name <subnet name>] - Subnet Name

Selects the subnets that match the given subnet name.

[-broadcast-domain <Broadcast Domain>] - Broadcast Domain

Selects the subnets that match the given broadcast domain name.

[-subnet <IP Address/Mask>] - Layer 3 Subnet

Selects the subnets that match the given address and mask.

[-gateway <IP Address>] - Gateway

Selects the subnets that match the given gateway address.

[-ip-ranges {<ipaddr>|<ipaddr>-<ipaddr>}] - IP Addresses or IP Address Ranges

Selects the subnets that match the given IP range.

[-total-count <integer>] - Total Address Count

Selects the subnets that match the given total address count.

[-used-count <integer>] - Used Address Count

Selects the subnets that match the given number of addresses allocated.

[-available-count <integer>] - Available Address Count

Selects the subnets that match the given number of addresses available.

Examples

The following example displays general information about the subnets.

```
cluster1::> network subnet show
IPspace: Default
     Subnet                      Broadcast                   Avail/
     Name      Subnet            Domain    Gateway           Total   Ranges
     --------- ---------------- --------- --------------- ---------
---------------
     s4        192.168.4.0/24   bd4       192.168.4.1        5/5
192.168.5.6-192.168.5.10
     s6        192.168.6.0/24   bd4       192.168.6.1        5/5
192.168.6.6-192.168.6.10
IPspace: ips1
     Subnet                      Broadcast                   Avail/
     Name      Subnet            Domain    Gateway           Total   Ranges
     --------- ---------------- --------- --------------- ---------
---------------
     s10       192.168.6.0/24   bd10      192.168.6.1        0/0    -
     3 entries were displayed.
```

network test-link commands

network test-link run-test

Test link bandwidth

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

Description

The `network test-link run-test` command runs a performance test between two nodes. The command requires a source node, Vserver, and destination address.

Before executing the `network test-link run-test` command, the network test-link start-server command must be run to start a server on the node hosting the destination LIF. After all tests to that node are complete the network test-link stop-server command must be run to stop the server.

The test results are stored non-persistently and can be viewed using the network test-link show command. Results include input parameters, the bandwidth achieved, and the date and time of the test.

Parameters

`-node {<nodename>|local}` - Node Name (privilege: advanced)

Use this parameter to specify the node which initiates the test.

`-vserver <vserver>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver to access the destination LIF. DC (Data Channel) Vserver option is available only in an ONTAP Select or Cloud Volumes ONTAP cluster. It is a special vserver that hosts LIFs that are used to mirror data aggregates to partner node.

`-destination <Remote InetAddress>` - Destination (privilege: advanced)

Use this parameter to specify the destination IP address.

Examples

The following example runs a test between the cluster LIFs, including the start and stop of the server side of the test:

```
cluster1::*> network test-link start-server -node node1

cluster1::*> network test-link run-test -node node2 -vserver Cluster
-destination 172.31.112.173
Node: node2
     Vserver: Cluster
 Destination: 172.31.112.173
Time of Test: 4/22/2016 15:33:18
        MB/s: 41.2678

cluster1::*> network test-link stop-server -node node1
cluster1::*> network test-link show
Node                Vserver             Destination       Time of Test
MB/s
-----------------   -----------------   ---------------
-------------------   ------------
node2               Cluster             172.31.112.173    4/22/2016
15:33:18         41.2678
```

Related Links

- network test-link start-server
- network test-link stop-server
- network test-link show

network test-link show

Display test results

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

Description

The `network test-link show` command displays the results of prior network test-link run-test commands.

The test results are stored non-persistently and can be viewed using the `network test-link show` command. Results include input parameters, the bandwidth achieved, and the date and time of the test.

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>|local} - Node (privilege: advanced)

Selects the nodes that match this parameter value. Use this parameter to display the test results specific to a node. By default, the test results across all nodes are shown.

-vserver <vserver> - Vserver (privilege: advanced)

Use this parameter to display the test results specific to a Vserver. Use DC (Data Channel) Vserver option only in an ONTAP Select or Cloud Volumes ONTAP cluster to show network performance of links hosting DC LIFs. DC Vserver is a special Vserver that hosts LIFs that are used to mirror data aggregates to partner node

[-destination <Remote InetAddress>] - Destination (privilege: advanced)

Use this parameter to display the test results associated with the specified destination.

[-timestamp <MM/DD/YYYY HH:MM:SS>] - Time of Test (privilege: advanced)

Use this parameter to display the test results with the specified timestamp.

[-bandwidth <double>] - MB/s (privilege: advanced)

Use this parameter to display the test results only matching the specified bandwidth.

Examples

The following example runs a test between the cluster LIFs twice and then demonstrates the show command results:

```
cluster1::*> network test-link run-test -node node2 -vserver Cluster
-destination 172.31.112.173
Node: node2
     Vserver: Cluster
 Destination: 172.31.112.173
Time of Test: 4/25/2016 10:37:52
        MB/s: 29.9946
cluster1::*> network test-link run-test -node node2 -vserver Cluster
-destination 172.31.112.173
Node: node2
     Vserver: Cluster
 Destination: 172.31.112.173
Time of Test: 4/25/2016 10:38:32
  MB/s: 39.8192
cluster1::*> network test-link show -node node2 -vserver Cluster
Node                Vserver             Destination       Time of Test
MB/s
-----------------   -----------------   ---------------
-------------------   ------------
node2               Cluster             172.31.112.173    4/25/2016
10:38:32         39.8192
```

Related Links

- network test-link run-test

network test-link start-server

Start server for bandwidth test

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

Description

The `network test-link start-server` command starts the server side of the `network test-link` test on the designated node.

Only one server at a time can be running for the `network test-link` command on a given node. If the `network test-link start-server` command is issued and a server is already running on the node, then the command is ignored, and the existing server continues to run.

The server started is listening on port 5201.

Parameters

`-node {<nodename>|local}` - Node Name (privilege: advanced)

Use this parameter to specify the node where the server is to be started.

Examples

The following example starts a server:

```
cluster1::*> network test-link start-server -node node1
```

network test-link stop-server

Stop server for bandwidth test

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

Description

The `network test-link stop-server` command stops the `network test-link` server running on the designated node.

Parameters

`-node {<nodename>|local}` - Node Name (privilege: advanced)

Use this parameter to specify the node where the server is to be stopped.

Examples

The following example starts a server and stops it:

```
cluster1::*> network test-link start-server -node node1

cluster1::*> network test-link stop-server -node node1
```

network tuning commands

network tuning icmp modify

Modify ICMP tuning options

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

Description

This command displays options which can be used to fine tune icmp protocol behavior.

Parameters

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

Sets this parameter to indicate on which node the ICMP tuning options are modified.

[-is-drop-redirect-enabled {true|false}] - Drop redirect ICMP

Sets this parameter to drop redirect ICMP message.

[-tx-icmp-limit <integer>] - Maximum number of ICMP packets sent per second

Sets the maximum number of ICMP messages including TCP RSTs can be sent per second.

[-redirect-timeout <integer>] - Maximum seconds for route redirect timeout

Sets this parameter to indicate the number of seconds after which the route is deleted. Value of zero means infinity. The default value is 300 seconds.

Examples

```
cluster1::> network tuning icmp modify -node node1 -is-drop-redirect
-enabled false
```

network tuning icmp show

Show ICMP tuning options

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

Description

This command displays the current state of the ICMP tuning options for the given node.

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] }

Displays all ICMP tuning options.

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

Specifies the node for which the ICMP tuning options are displayed.

[-is-drop-redirect-enabled {true|false}] - Drop redirect ICMP

Displays all entries that match the "is-drop-redirect-enabled" value.

[-tx-icmp-limit <integer>] - Maximum number of ICMP packets sent per second

Displays all entries that match the "tx-icmp-limit" value.

[-redirect-timeout <integer>] - Maximum seconds for route redirect timeout

Displays all the entries that match the "redirect-timeout" value.

Examples

```
cluster1::> network tuning icmp show
Drop Redirect Maximum ICMP      Redirect Timeout
Node      ICMP          Sends per Second  in Seconds
--------- ------------- ----------------  ----------------
node1
          true          100               300
```

network tuning icmp6 modify

Modify ICMPv6 tuning options

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

Description

This command displays options which can be used to fine tune icmpv6 protocol behavior.

Parameters

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

Sets this parameter to indicate on which node the ICMPv6 tuning options are modified.

`[-is-v6-redirect-accepted {true|false}]` - Accept redirects via ICMPv6

Sets this parameter to indicate whether or not redirect ICMPv6 messages are accepted.

`[-redirect-v6-timeout <integer>]` - Maximum seconds for route redirect timeout

Sets this parameter to indicate the number of seconds after which the route is deleted. Value of zero means infinity. The default value is 300 seconds.

`[-tx-icmp6-err-limit <integer>]` - Maximum number of ICMPv6 error messages sent per second

Sets the maximum number of ICMPv6 error messages that can be sent per second.

Examples

```
cluster1::> network tuning icmp6 modify -node node1 -is-v6-redirect
-accepted false
```

network tuning icmp6 show

Show ICMPv6 tuning options

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

Description

This command displays the current state of the ICMPv6 tuning options for the given node.

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] }

Displays all ICMPv6 tuning options.

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

Specifies the node for which the ICMPv6 tuning options are displayed.

[-is-v6-redirect-accepted {true|false}] - Accept redirects via ICMPv6

Displays all entries that match the "is-v6-redirect-accepted" value.

[-redirect-v6-timeout <integer>] - Maximum seconds for route redirect timeout

Displays all the entries that match the "redirect-v6-timeout" value.

[-tx-icmp6-err-limit <integer>] - Maximum number of ICMPv6 error messages sent per second

Displays all entries that match the "tx-icmp6-err-limit" value.

Examples

```
cluster1::> network tuning icmp6 show
Accept Redirect Maximum ICMPv6 Error Redirect Timeout
Node      ICMPv6          Sends per Second     in Seconds
--------- --------------- -------------------- ----------------
node1
          true            100                  300
```

network tuning tcp modify

Modify TCP tuning options

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

Description

This commands sets TCP tuning options on the node.

Parameters

-node {<nodename>|local} - Node

Indicates on which node the TCP tuning options will be modified.

[-is-path-mtu-discovery-enabled {true|false}] - Path MTU discovery enabled

Enables path MTU discovery feature.

[-is-rfc3465-enabled {true|false}] - RFC3465 enabled

Enables the rfc3465 feature.

[-max-cwnd-increment <integer>] - Maximum congestion window segments incrementation

Sets the maximum congestion window increment segments during slow start.

[-is-rfc3390-enabled {true|false}] - RFC3390 enabled

Enables the rfc3390 feature.

[-is-sack-enabled {true|false}] - SACK support enabled

Enables the selective ACK feature.

Examples

```
cluster1::> network tuning tcp modify -node node1 -is-path-mtu-discovery
-enabled false
```

network tuning tcp show

Show TCP tuning options

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

Description

This command displays the current state of the TCP tuning options for the given node.

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] }

Displays all TCP tuning options.

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

Specifies the node for which the TCP tuning options will be displayed.

[-is-path-mtu-discovery-enabled {true|false}] - Path MTU discovery enabled

Displays all entries that match the "is-path-mtu-discovery-enabled" value.

[-is-rfc3465-enabled {true|false}] - RFC3465 enabled

Displays all entries that match the "is-rfc3465-enabled" value.

[-max-cwnd-increment <integer>] - Maximum congestion window segments incrementation

Displays all entries that match the "max-cwnd-increment" value.

[-is-rfc3390-enabled {true|false}] - RFC3390 enabled

Displays all entries that match the "is-rfc3390-enabled" value.

[-is-sack-enabled {true|false}] - SACK support enabled

Displays all entries that match the "is-sack-enabled" value.

Examples

```
cluster1::> network tuning tcp show
          Path MTU           Maximum                    Selective
Node      Discovery  RFC3465 Congestion Window RFC3390 Ack
          Enabled    Enabled Incrementation    Enabled Enabled
--------- ---------  ------- ----------------- ------- --------
node1
          true       true    2                 true    true
```

protection-type commands

protection-type show

Display the supported protection types and available RPOs

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

Description

This command displays the protection types available for application provisioning.

Parameters

`[-fields <fieldname>,…]`

Specifies fields that you want included in the output. You can use `-fields ?` to display the available fields.

| `[-instance ]` (privilege: advanced) }

Specifies the display of all available fields for each selected protection type.

`[-vserver <vserver name>]` - Vserver (privilege: advanced)

Selects the protection types of Vservers that match the parameter value.

`[-protection-type {local|remote}]` - Protection Type (privilege: advanced)

Selects the protection types that match the parameter value.

`[-rpo-list <text>,…]` - List of available RPOs (privilege: advanced)

Selects the protection types whose list of available RPOs matches the parameter value.

`[-rpo-list-description <text>,…]` - List of descriptions of available RPOs (privilege: advanced)

Selects the protection types whose list of description of available RPOs matches the parameter value.

`[-description <text>]` - Description of Protection Type (privilege: advanced)

Selects the protection types with a description that matches the parameter value.

Examples

The following example displays the protection types and the associated available RPOs for all Vservers in the cluster.

```
cluster1::*> protection-type show
Vserver Protection Type RPO List
------- --------------- --------------------
vs1
        local           hourly, none
  remote          none
vs2
        local           hourly, none
  remote          none, zero
```

qos commands

qos adaptive-policy-group commands

qos adaptive-policy-group create

Create an adaptive policy group

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

Description

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

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

Parameters

`-policy-group <text>` - Name

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

`-vserver <vserver name>` - Vserver

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Examples

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

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

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

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

Related Links

- qos adaptive-policy-group rename

qos adaptive-policy-group delete

Delete an adaptive policy group

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

Description

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

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

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

Parameters

`-policy-group <text>` - Name

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

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

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

Examples

```
The following example deletes "p1" adaptive policy group:
```

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

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

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

qos adaptive-policy-group modify

Modify an adaptive policy group

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

Description

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

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

Parameters

`-policy-group <text>` - Name

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Examples

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

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

Related Links

- qos adaptive-policy-group rename

qos adaptive-policy-group rename

Rename an adaptive policy group

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

Description

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

Parameters

`-policy-group <text>` - Name

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

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

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

Examples

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

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

qos adaptive-policy-group show

Display a list of adaptive policy groups

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

Description

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

Parameters

{ [-fields <fieldname>,…]

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

| [-instance] }

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

[-policy-group <text>] - Name

Selects the adaptive policy groups that match this parameter value.

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

[-vserver <vserver name>] - Vserver

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

[-uuid <UUID>] - Uuid

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

[-pgid <integer>] - ID

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Examples

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

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

qos policy-group commands

qos policy-group create

Create a policy group

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

Description

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

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

Parameters

`-policy-group <text>` - Policy Group Name

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

`-vserver <vserver name>` - Vserver

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

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

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

internally treated as infinity.

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

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

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

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

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

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

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

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

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

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

Examples

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

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

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

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

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

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

Related Links

- qos policy-group rename

qos policy-group delete

Delete an existing QoS Policy Group

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

Description

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

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

Parameters

`-policy-group <text>` - Policy Group Name

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

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

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

Examples

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

Deletes the "p1" policy group.

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

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

qos policy-group modify

Modify a policy group

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

Description

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

Parameters

`-policy-group <text>` - Policy Group Name

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

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

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

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

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

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

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

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

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

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

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

Examples

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

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

qos policy-group rename

Rename a policy group

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

Description

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

Parameters

`-policy-group <text>` - Policy Group Name

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

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

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

Examples

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

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

qos policy-group show

Display a list of policy groups

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

Description

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

Parameters

{ `[-fields <fieldname>,…]`

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

| `[-instance ]` }

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

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

Selects the policy groups that match this parameter value

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

`[-vserver <vserver name>]` - Vserver

Selects the policy groups that match this parameter value

`[-uuid <UUID>]` - Uuid

Selects the policy groups that match this parameter value

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

Selects the policy groups that match this parameter value

[-pgid <integer>] - Policy Group ID

Selects the policy groups that match this parameter value

This uniquely identifies the policy group

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

Selects the policy groups that match this parameter value

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

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

Selects the policy groups that match this parameter value

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

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

Selects the policy groups that match this parameter value.

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

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

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

Selects the policy groups that match this parameter value.

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

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

Selects the policy groups that match this parameter value.

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

Examples

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

The example above displays all policy groups on the cluster.

qos settings commands

qos settings throughput-floors-v2

Enable/Disable throughput floors v2 on AFF

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

Description

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

Parameters

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

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

Examples

The following example disables floors v2 on AFF.

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

qos settings cache modify

Modify the cache policy

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

Description

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

Parameters

`-cache-setting <text>` - Cache Policy Name

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

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

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

Examples

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

Sets caching-policy random_read_write-random_write as default.

qos settings cache show

Display list of cache policies

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

Description

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

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

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

Parameters

{ [-fields <fieldname>,…]

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

| [-instance] }

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

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

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

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

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

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

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

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

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

Examples

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

Shows QoS settings for the caching policies.

qos statistics commands

qos statistics characteristics show

Display QoS policy group characterization

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

Description

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

The command displays the following data:

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

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

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

Parameters

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

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

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

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

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

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

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

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

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

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

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

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

Examples

```
cluster1::> qos statistics characteristics show -iterations 100 -rows 4
Policy Group             IOPS      Throughput Request size  Read
Concurrency
-------------------- -------- --------------- ------------  ----
-----------
-total-                    31      304.00KB/s       10041B    0%
16
_System-Best-Effort        15           0KB/s           0B    0%
0
vol1                       11       44.00KB/s        4096B    0%
40
vol2                        4      256.00KB/s       65536B    0%
14
vs1vol0                     1        4.00KB/s        4096B    0%
4
-total-                    37      808.00KB/s       22361B    2%
3
_System-Best-Effort        15           0KB/s           0B    0%
0
vol2                       12      768.00KB/s       65536B    0%
9
vs1vol0                     8       32.00KB/s        4096B   12%
1
vol1                        2        8.00KB/s        4096B    0%
1
```

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

```
cluster1::> qos statistics characteristics show -iterations 100 -policy
-group pg1
Policy Group             IOPS      Throughput Request size Read
Concurrency
-------------------- -------- --------------- ------------ ----
-----------
-total-                   293        3.02MB/s       10783B  54%
0
pg1                       118      470.67KB/s        4096B 100%
0
-total-                   181      478.14KB/s        2700B  65%
0
pg1                       117      469.33KB/s        4096B 100%
0
-total-                   226      525.78KB/s        2382B  60%
1
pg1                       110      440.00KB/s        4096B 100%
1
-total-                   233        1.67MB/s        7527B  49%
1
pg1                       112      446.67KB/s        4096B 100%
1
```

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

qos statistics latency show

Display latency breakdown data per QoS policy group

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

Description

The `qos statistics latency show` command displays the average latencies for QoS policy groups across the various 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 SnapMirror Synchronous (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
-------------------- ---------- ---------- ---------- ----------
---------- ---------- ---------- ---------- ---------- ----------
---------- ----------
-total-                110.35ms   110.02ms        0ms   327.00us
```

```
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       Data
Disk    QoS Max    QoS Min     NVRAM       Cloud  FlexCache   SM Sync
VA
-------------------- ---------- ---------- ---------- ----------
---------- ---------- ---------- ---------- ---------- ----------
---------- ----------
-total-                  5.88ms   308.00us        0ms   434.00us
5.14ms        0ms        0ms         0ms       0ms        0ms        0ms
0ms
pg1                      5.88ms   308.00us        0ms   434.00us
5.14ms        0ms        0ms         0ms       0ms        0ms        0ms
0ms
-total-                  4.17ms   280.00us        0ms   477.00us
3.42ms        0ms        0ms         0ms       0ms        0ms        0ms
0ms
pg1                      4.17ms   280.00us        0ms   477.00us
3.42ms        0ms        0ms         0ms       0ms        0ms        0ms
0ms
-total-                  4.43ms   274.00us        0ms   656.00us
3.50ms        0ms        0ms         0ms       0ms        0ms        0ms
0ms
pg1                      4.43ms   274.00us        0ms   656.00us
3.50ms        0ms        0ms         0ms       0ms        0ms        0ms
0ms
-total-                  4.89ms   276.00us        0ms   699.00us
3.92ms        0ms        0ms         0ms       0ms        0ms        0ms
0ms
pg1                      4.89ms   276.00us        0ms   699.00us
3.92ms        0ms        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
-policy-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
-policy-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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

Examples

```
cluster1::> qos statistics volume characteristics show -iterations 100
-rows 3
Workload            ID     IOPS       Throughput Request size Read
Concurrency
--------------- ------ -------- ---------------- ------------ ----
-----------
-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_vol0 -iterations 100
Workload            ID     IOPS       Throughput Request Size Read
Concurrency
--------------- ------ -------- ---------------- ------------ ----
-----------
-total-              -     1567       783.33KB/s        512Kb  90%
2
vs0_vol0-wid1..  15658      785       392.33KB/s        512Kb  89%
1
-total-              -     1521       760.50KB/s        512Kb  90%
1
vs0_vol0-wid1..  15658      982       491.17KB/s        512Kb  90%
0
-total-              -     1482       741.00KB/s        512Kb  89%
0
vs0_vol0-wid1..  15658      945       472.50KB/s        512Kb  90%
0
-total-              -     1482       741.00KB/s        512Kb  89%
0
vs0_vol0-wid1..  15658      945       472.50KB/s        512Kb  90%
0
-total-              -     1702       850.83KB/s        512Kb  90%
0
vs0_vol0-wid1..  15658     1018       509.00KB/s        512Kb  90%
0
```

The example above displays characteristics for volume *`vs0_vol0`* 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 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 SnapMirror Synchronous (SM Sync) operations
- Latency observed per I/O operation for Volume Activation (VA) operations
- Latency observed per I/O operation for Anti Virus Scan (AVSCAN) 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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

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
AVSCAN
--------------- ------ --------   -------- --------   -------- --------
-------- ---------- ---------- ---------- ---------- ---------- ----------
----------
-total-                110.35ms   110.02ms      0ms   327.00us      0ms
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
0ms
vol1              1234 117.76ms   117.56ms      0ms   191.00us      0ms
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
0ms
-total-             -   38.89ms    38.63ms      0ms   256.00us      0ms
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
0ms
vol1              1234  27.28ms    27.03ms      0ms   253.00us      0ms
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
0ms
-total-              - 409.81ms   409.65ms      0ms   169.00us      0ms
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
0ms
vol2               999 407.88ms   407.66ms      0ms   219.00us      0ms
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
0ms
```

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

```
cluster1::> qos statistics volume latency show -vserver vs0 -volume
vs0_vol0 -iterations 100
Workload            ID    Latency    Network    Cluster       Data
Disk    Qos Max    Qos Min      NVRAM      Cloud  FlexCache    SM Sync
VA     AVSCAN
--------------- ------ ---------- ---------- ---------- ----------
---------- ---------- ---------- ---------- ---------- ----------
---------- ---------- ----------
-total-              -   455.00us   158.00us        0ms   297.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
vs0_vol0-wid1..  15658   428.00us   155.00us        0ms   273.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
-total-              -   337.00us   130.00us        0ms   207.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
vs0_vol0-wid1..  15658   316.00us   128.00us        0ms   188.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
-total-              -   464.00us   132.00us        0ms   332.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
vs0_vol0-wid1..  15658   471.00us   130.00us        0ms   341.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
-total-              -   321.00us   138.00us        0ms   183.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
vs0_vol0-wid1..  15658   302.00us   137.00us        0ms   165.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
-total-              -   418.00us   142.00us        0ms   276.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
vs0_vol0-wid1..  15658   424.00us   143.00us        0ms   281.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
```

The example above displays latencies for volume *vs0_vol0* 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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

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_vol0 -iterations 100
Workload            ID     IOPS       Throughput    Latency
--------------- ------ -------- ---------------- ----------
-total-              -     1278       639.17KB/s   404.00us
vs0_vol0-wid1..  15658      526       263.17KB/s   436.00us
-total-              -     1315       657.33KB/s    86.00us
vs0_vol0-wid1..  15658      528       264.17KB/s    88.00us
-total-              -     1220       609.83KB/s   418.00us
vs0_vol0-wid1..  15658      515       257.33KB/s   531.00us
-total-              -     1202       600.83KB/s   815.00us
vs0_vol0-wid1..  15658      519       259.67KB/s   924.00us
-total-              -     1240       620.17KB/s   311.00us
vs0_vol0-wid1..  15658      525       262.50KB/s   297.00us
```

The example above displays the system performance for volume *vs0_vol0* 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-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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

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_vol0-wid-..   212    2%
    -total- (100%)      -    8%
    vs0vol1-wid-102   102    5%
    vs0vol2-wid-121   121    2%
    vs2_vol0-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_vol1 -iterations 100
    Workload           ID   CPU
    --------------- ----- -----
    -total- (100%)      -    2%
    vs0_vol1-wid7..  7916    2%
    -total- (100%)      -    2%
    vs0_vol1-wid7..  7916    2%
    -total- (100%)      -    1%
    vs0_vol1-wid7..  7916    1%
    -total- (100%)      -    2%
    vs0_vol1-wid7..  7916    1%
    -total- (100%)      -    2%
    vs0_vol1-wid7..  7916    2%
```

The example above displays total CPU utilization for volume *vs0_vol1* 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-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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

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_vol0 -iterations 100
            Workload            ID   Disk Number of HDD Disks
            --------------- ------ ------ -------------------
            -total-              -     5%                  10
            vs0_vol0-wid1..  15658     1%                   6
            -total-              -     5%                  10
            vs0_vol0-wid1..  15658     1%                   6
            -total-              -     6%                  10
            vs0_vol0-wid1..  15658     2%                   6
            -total-              -     6%                  10
            vs0_vol0-wid1..  15658     2%                   6
            -total-              -     6%                  10
            vs0_vol0-wid1..  15658     2%                   6
```

The example above displays total disk utilization for volume *vs0_vol0* 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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

Examples

```
cluster1::> qos statistics workload characteristics show -iterations 100
-rows 4
Workload            ID     IOPS       Throughput Request size Read
Concurrency
--------------- ------ -------- ---------------- ------------ ----
-----------
-total-              -       68       176.00KB/s        2650B   7%
8
vs1vol0-wid102     102       24        96.00KB/s        4096B  20%
13
_Scan_Besteff..    101       23            0KB/s           0B   0%
0
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
_Scan_Besteff..    101       25            0KB/s           0B   0%
0
-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
_Scan_Besteff..    101       20            0KB/s           0B   0%
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
    file-test1_a-..   6437       34       136.00KB/s        4096B 100%
0
    file-test1_c-..   5078       33       133.33KB/s        4096B 100%
0
    -total-              -      310         3.09MB/s       10428B  55%
1
    file-test1_a-..   6437       36       142.67KB/s        4096B 100%
0
    file-test1_b-..   9492       35       138.67KB/s        4096B 100%
0
    -total-              -      192       575.71KB/s        3075B  71%
1
    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
-workload-id 9492
    Workload            ID     IOPS       Throughput Request size Read
Concurrency
    --------------- ------ -------- ---------------- ------------ ----
-----------
    -total-              -      737         2.14MB/s        3045B  79%
1
    file-test1_b-..   9492      265      1058.67KB/s        4096B 100%
0
    -total-              -      717         4.26MB/s        6235B  80%
1
    file-test1_b-..   9492      272      1086.67KB/s        4096B 100%
1
    -total-              -      623         2.50MB/s        4202B  86%
0
    file-test1_b-..   9492      263      1050.67KB/s        4096B 100%
0
    -total-              -      595         2.11MB/s        3712B  89%
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 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 SnapMirror Synchronous (SM Sync) operations
- Latency observed per I/O operation for Volume Activation (VA) operations
- Latency observed per I/O operation for Anti Virus Scan (AVSCAN) 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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

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
AVSCAN
--------------- ------ --------   -------- --------   -------- --------
-------- ---------- ---------- ---------- ---------- ---------- ----------
----------
-total-                110.35ms   110.02ms      0ms   327.00us      0ms
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
0ms
vol1              1234 117.76ms   117.56ms      0ms   191.00us      0ms
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
0ms
-total-             -   38.89ms    38.63ms      0ms   256.00us      0ms
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
0ms
vol1              1234  27.28ms    27.03ms      0ms   253.00us      0ms
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
0ms
-total-              - 409.81ms   409.65ms      0ms   169.00us      0ms
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
0ms
vol2               999 407.88ms   407.66ms      0ms   219.00us      0ms
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
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     AVSCAN
--------------- ------ ---------- ---------- ---------- ----------
---------- ---------- ---------- ---------- ---------- ----------
---------- ---------- ----------
-total-              -     4.80ms   287.00us        0ms   427.00us
4.08ms        0ms        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        0ms
file-test1_a-..   6437     8.22ms   262.00us        0ms   424.00us
7.53ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
-total-              -     4.20ms   296.00us        0ms   421.00us
3.48ms        0ms        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        0ms
file-test1_a-..   6437     6.70ms   297.00us        0ms   464.00us
5.94ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
-total-              -     5.90ms   303.00us        0ms     1.71ms
3.88ms        0ms        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        0ms
file-test1_a-..   6437     9.48ms   250.00us        0ms     2.30ms
6.93ms        0ms        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.

```
cluster1::> qos statistics workload latency show -iterations 100 -workload
-id 9492
   Workload            ID    Latency    Network    Cluster       Data
Disk    Qos Max    Qos Min      NVRAM      Cloud  FlexCache    SM Sync
VA     AVSCAN
   --------------- ------ ---------- ---------- ---------- ----------
---------- ---------- ---------- ---------- ---------- ----------
---------- ---------- ----------
   -total-              -   443.00us   273.00us        0ms   170.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
   file-test1_b-..   9492   440.00us   272.00us        0ms   168.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
   -total-              -   577.00us   313.00us        0ms   264.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
   file-test1_b-..   9492   607.00us   316.00us        0ms   291.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
   -total-              -   475.00us   291.00us        0ms   184.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
   file-test1_b-..   9492   476.00us   293.00us        0ms   183.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
   -total-              -   628.00us   284.00us        0ms   344.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
   file-test1_b-..   9492   591.00us   281.00us        0ms   310.00us
0ms        0ms        0ms        0ms        0ms        0ms        0ms
0ms        0ms
```

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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

Examples

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

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

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

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

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

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

qos statistics workload resource cpu show

Display CPU resource utilization data per QoS workload

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

Description

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

The command displays the following data:

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

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

Parameters

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

Selects the QOS workloads that match this parameter value.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

If the parameter is specified and if the value is true, it will display data for FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

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_vol0-wid-..   212    2%
    -total- (100%)      -    8%
    vs0-wid-101       102    5%
    file-bigvmdk-..   121    2%
    vs2_vol0-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 FlexVol volumes and FlexGroup constituents. Otherwise it will display data for FlexVol volumes and FlexGroup volumes.

Examples

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

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

```
cluster1::> qos statistics workload resource disk show -node local
-iterations 100 -rows 2 -policy-group pg1
Workload            ID  Disk Number of HDD Disks
--------------- ------ ----- -------------------
-total-              -    3%                  10
file-test1_a-..   6437    6%                   6
file-test1-wi..   7872    6%                   6
-total-              -    3%                  10
file-test1_a-..   6437    5%                   6
file-test1-wi..   7872    5%                   6
-total-              -    3%                  10
file-test1_a-..   6437    6%                   6
file-test1-wi..   7872    6%                   6
```

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

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

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

qos workload commands

qos workload delete

Delete workload

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

Description

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

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

Parameters

`-workload <text>` - Workload Name

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

Examples

```
cluster1::> qos workload delete workload1
```

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

qos workload show

Display a list of workloads

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

Description

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

Parameters

{ [-fields <fieldname>,…]

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

| [-instance] }

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

[-workload <text>] - Workload Name

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

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

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

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

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

[-wid <integer>] - Workload ID

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

[-category <text>] - Workload Category

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

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

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

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

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

[-vserver <vserver name>] - Vserver

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

[-volume <volume name>] - Volume

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

[-qtree <qtree name>] - Qtree Name

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

[-lun <text>] - LUN Name

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

[-file <text>] - File Path

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

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

Selects the workloads that match this parameter value

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

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

Selects the workloads that match this parameter value

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

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

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

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

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

Examples

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

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

san commands

san config commands

san config show

Show SAN configuration options for the cluster

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

Description

The command displays cluster SAN configuration options.

Examples

The following example displays the cluster SAN configuration.

```
cluster1::> san config show
        All SAN Array: true
```

security commands

security snmpusers

Show SNMP users

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

Description

The `security snmpusers` displays the following information about SNMP users:

- User name
- Authentication method
- Hexadecimal engine ID
- Authentication protocol
- Privacy protocol
- Security group

Parameters

{ [-fields <fieldname>,…]

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

| [-instance] }

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

[-vserver <Vserver Name>] - Vserver

If this parameter is specified, the command displays information only about the SNMP user or users that belong to the specified Vserver.

[-username <text>] - User Name

If this parameter is specified, the command displays information only about the SNMP user with the specified user name.

[-authmethod <text>] - Authentication Method

If this parameter is specified, the command displays information only about the SNMP user or users that use the specified authentication method. Possible values include the following:

- community-SNMP community strings
- usm-SNMP user security model

[-remote-switch-ipaddress <IP Address>] - Remote Switch IP Address

If this parameter is specified, the command displays information only about the remote SNMP user or users that belong to the specified remote switch.

[-engineid <Hex String>] - Engine Id

If this parameter is specified, the command displays information only about the SNMP user or users that use the specified engine ID, specified in hexadecimal format.

[-authprotocol <text>] - Authentication Protocol

If this parameter is specified, the command displays information only about the SNMP user or users that use the specified authentication protocol.

[-privprotocol <text>] - Privacy Protocol

If this parameter is specified, the command displays information only about the SNMP user or users that use the specified privacy protocol.

[-securitygroup <text>] - Security Group

If this parameter is specified, the command displays information only about the SNMP user or users that belong to the specified security group.

Examples

The following example displays information about all SNMP users:

```
cluster1::> security snmpusers
                                                   Protocols Security  Remote
Vserver  UserName   AuthMethod EngineId         Auth Priv Group     Switch
IP
-------- ---------- ---------- ---------------- ---- ---- ---------
-----------
cluster1 comm1      community  8000031504312d38302d313233343536     -
                                                   -    -    readwrite
cluster1 private    community  8000031504312d38302d313233343536      -
                                                   -    -    readwrite
cluster1 snmpuser1  usm        80000634b21000000533296869
                                                   -    -    readwrite
172.2.20.91
vs1      snmpuser2  community  8000031504312d38302d31323334353632
                                                   -    -    readwrite -
vs1      snmpuser3  usm        8000031504312d38302d31323334353632
                                                   -    -    readwrite -
```

security anti-ransomware commands

security anti-ransomware show

Display node Anti-ransomware Version

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

Description

The `security anti-ransomware show` command displays the Anti-ransomware version of a node.

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>|local}] - Node Name

If this parameter is specified, the command displays the Anti-ransomware version of specified node.

[-version <text>] - Anti-Ransomware Version

If this parameter is specified, the command displays node with the specified Anti-ransomware version.

Examples

The following example shows a sample output for this command:

```
C1_cluster::*> security anti-ransomware show-version
Node                 Version
-------------------- ----------------
sti221-vsim-sr049g   1.0
sti221-vsim-sr049h   1.0
2 entries were displayed.
```

security anti-ransomware update-package-from-uri

Install the new Anti-ransomware package in the cluster

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

Description

The `security anti-ransomware update-package-from-uri` command downloads a new Anti-ransomware package to replace the version installed in all nodes of the cluster.

Use only Anti-ransomware package provided by NetApp support for ONTAP with this command."

To update the Anti-ransomware package, perform the following steps:

- Download the required Anti-ransomware package file from the NetApp support website.
- Place the downloaded package file (.tgz) on a local web server that the cluster can access without a password.
- Execute the `security anti-ransomware update-package-from-uri` command, passing the

Universal Resource Identifier (URI) of the file as a parameter.

> "The command only needs to be executed once for the entire cluster. It will then download and install the package on each node."

Parameters

`-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - URI of Anti-ransomware package

URI of the new Anti-ransomware package.

Examples

The following example installs a new version of the Anti-ransomware to the cluster:

```
cluster1::> security anti-ransomware update-package-from-uri
http://www.example.com/auDev__ontap_arpml_package_v1.0.tgz__440ae2e4-fd8f-
4225-9bee-94e2da3f9d8d.tgz
```

security anti-ransomware volume disable

Disable anti-ransomware on a volume

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

Description

The `security anti-ransomware volume disable` command disables anti-ransomware monitoring on a volume.

Parameters

`-vserver <vserver name>` - Vserver Name

Anti-ransomware monitoring is disabled on volumes that match the values for the Vserver and volume parameters. If only one data Vserver exists, you do not need to specify this parameter.

`-volume <volume name>` - Volume Name

Anti-ransomware monitoring is disabled on volumes matching the parameter value.

Examples

security anti-ransomware volume dry-run

Dry-run anti-ransomware on a volume

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

Description

The `security anti-ransomware volume dry-run` command starts anti-ransomware monitoring in the evaluation mode on a volume.

Parameters

`-vserver <vserver name>` - Vserver Name

Anti-ransomware monitoring is enabled in the evaluation mode on volumes that match the values for the Vserver and volume parameters. If only one data Vserver exists, you do not need to specify this parameter.

`-volume <volume name>` - Volume Name

Anti-ransomware monitoring is enabled in the evaluation mode on volumes matching the parameter value.

Examples

security anti-ransomware volume enable

Enable anti-ransomware on a volume

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

Description

The `security anti-ransomware volume enable` command enables anti-ransomware monitoring on a volume.

Parameters

`-vserver <vserver name>` - Vserver Name

Anti-ransomware monitoring is enabled on volumes that match the values for the Vserver and volume parameters. If only one data Vserver exists, you do not need to specify this parameter.

`-volume <volume name>` - Volume Name

Anti-ransomware monitoring is enabled on volumes matching the parameter value.

Examples

security anti-ransomware volume pause

Pause anti-ransomware on a volume

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

Description

The `security anti-ransomware volume pause` command pauses Anti-ransomware monitoring on a volume.

Parameters

`-vserver <vserver name>` - Vserver Name

Anti-ransomware monitoring is paused in the evaluation mode on volumes that match the values for the Vserver and volume parameters. If only one data Vserver exists, you do not need to specify this parameter.

`-volume <volume name>` - Volume Name

Anti-ransomware monitoring is paused on volumes matching the parameter value.

Examples

security anti-ransomware volume resume

Resume anti-ransomware on a volume

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

Description

The `security anti-ransomware volume resume` command resumes Anti-ransomware monitoring on a volume.

Parameters

`-vserver <vserver name>` - Vserver Name

Anti-ransomware monitoring is resumed on volumes that match the values for the Vserver and volume parameters. If only one data Vserver exists, you do not need to specify this parameter.

`-volume <volume name>` - Volume Name

Anti-ransomware monitoring is resumed on volumes matching the parameter value.

Examples

security anti-ransomware volume show

Show anti-ransomware related information of volumes

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

Description

The `security anti-ransomware volume show` command displays information related to Anti-ransomware on the volumes in the cluster. The following information is displayed:

- Vserver Name: The Vserver on which the volume is located.
- Volume Name: The volume name
- State: The Anti-ransomware state of the volume. The possible values are *`disabled`* , *`enabled`* , *`dry-run`* , *`dry-run-paused`* , *`enable-paused`* and *`disable-in-progress`* .

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.

| [-attack]

If this parameter is specified, ransomware attack details are displayed.

| [-instance] }

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

[-vserver <vserver name>] - Vserver Name

If this parameter and the `-volume` parameter are specified, the command displays detailed information related to Anti-ransomware about the specified volume. If this parameter is specified by itself, the command displays information related to the Anti-ransomware about all volumes on the specified Vserver.

[-volume <volume name>] - Volume Name

If this parameter and the `-vserver` parameter are specified, the command displays detailed information related to Anti-ransomware about the specified volume. If this parameter is specified by itself, the command displays information related to the Anti-ransomware about all volumes matching the specified name.

[-state {disabled|enabled|dry-run|paused|dry-run-paused|enable-paused|disable-in-progress}] - State

If this parameter is specified, the command displays information only about the volume or volumes that have the specified Anti-ransomware state. The possible values are *disabled* , *enabled* , *dry-run* , *dry-run-paused* , *enable-paused* and *disable-in-progress* . The possible states are:

- disabled - Anti-ransomware is disabled on the volume.
- enabled - Anti-ransomware is enabled on the volume.
- dry-run - Anti-ransomware is enabled in the dry-run or evaluation mode on the volume.
- dry-run-paused - Anti-ransomware is paused from dry-run or evaluation mode on the volume.
- enable-paused - Anti-ransomware is paused on the volume.
- disable-in-progress - Anti-ransomware disable work is in progress on the volume.

[-dry-run-start-time <MM/DD/YYYY HH:MM:SS>] - Dry Run Start Time

If this parameter is specified, the command displays the dry run start time of the volumes that have the state dry-run or dry-run-paused.

[-attack-probability {none|low|moderate|high}] - Attack Probability

If this parameter is specified, the command displays information only about the volumes that have the specified probability. The possible values are *none* ,*low* , *moderate* , and *high* .

- none - No data is suspected for ransomware activity.
- low - Small amount data is suspected for ransomware activity.
- moderate - Moderate amount of data is suspected for ransomware activity.
- high - Large amount data is suspected for ransomware activity.

[-attack-timeline <MM/DD/YYYY HH:MM:SS>,…] - Attack Timeline

If this parameter is specified, the command displays information only about the volumes that have the specified attack-timeline.

[-no-of-attacks <integer>] - Number of Attacks

This provides the number of ransomware attacks observed.

Examples

The following example shows a sample output for this command:

```
cluster1::> security anti-ransomware volume show

Vserver    Volume       State
---------- ------------ -------
vs1        vol1         enabled
```

security anti-ransomware volume attack clear-suspect

Clear suspect record

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

Description

The `anti-ransomware volume attack clear-suspect` command removes the specified files from suspect files report. When no optional parameters are provided, the suspect report file is cleared. If the attack is marked as a true positive, by setting the `false-positive` parameter to *`false`* , the snapshot is retained for 7 days. If marked as a false positive, by setting the `false-positive` parameter to *`true`* , the snapshot is retained for 24 hours.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume on which anti-ransomware feature is enabled.

{ [-sequence-number <integer>] - Sequence Number

This optionally specifies the sequence number of the suspect file obtained from generated report.

| [-extensions <text>,…] - File Extensions

This optionally specifies the extensions of ransomware attacked files that needs to be cleared from attack report.

| [-start-time <MM/DD/YYYY HH:MM:SS>] - Start Time

This optionally specifies the lower bound of the time to clear a suspect record. Any suspect record with time greater than or equal to start-time is cleared.

[-end-time <MM/DD/YYYY HH:MM:SS>] - End Time }

This optionally specifies upper bound of the time to clear a suspect record. Any suspect record with time less than or equal to end-time is cleared.

-false-positive {true|false} - False Positive?

This indicates whether the suspect record of specific extensions, time range, and so on, are to be considered a false positive.

Examples

The following example shows a sample output for clearing all the suspects observed with timestamp in the start-time and end-time range, and with given extension.

```
clus1::> security anti-ransomware volume attack clear-suspect -volume
testvol -start-time "4/14/2021 04:16:48" -end-time "4/14/2021 06:16:50"
5 suspect records cleared.

The following examples shows output when given sequence-number is not
present.
clus1::*> security anti-ransomware volume attack clear-suspect -volume
testvol -sequence-number 1000

Error: command failed: No suspect records found.
```

security anti-ransomware volume attack generate-report

Generates Report File of the Suspected Attack on the Volume

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

Description

The `anti-ransomware volume attack generate-report` command copies the report file to the given path. The report file is not an exhaustive list of all suspected files during the attack, since the maximum number of entries allowed per file extension is 20.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume on which anti-ransomware feature is enabled.

-dest-path <sub dir> - Destination path under the volume to copy the report file

This parameter specifies the path where requested file is to be copied.

Examples

The following example displays command output:

```
node::*> security anti-ransomware volume attack generate-report -volume
vol1 -dest-path vs1:vol1/
Report "report_file_vs1vol1_30-03-2021_16-11-38" available at path
"vs1:vol1/".
```

security anti-ransomware volume attack-detection-parameters modify

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

Description

The `security anti-ransomware volume attack-detection-parameters modify` command can be used to modify the attack detection parameters of an anti-ransomware enabled volume.

Parameters

`-vserver <Vserver Name>` - Vserver Name

This parameter specifies the Vserver of the anti-ransomware enabled volume.

`-volume <volume name>` - Volume Name

This parameter specifies the anti-ransomware enabled volume for which the attack detection parameters need to be modified.

`[-based-on-high-entropy-data-rate {true|false}]` - High Entropy Data Rate at Volume Level

This parameter indicates whether ransomware detection is based on a high entropy data rate at the volume level. Ransomware detection is also done based on high entropy data rate at the file level and this method of detection is always enabled and has no dependency on this parameter.

`[-based-on-never-seen-before-file-extension {true|false}]` - Never Seen before File Extension

This parameter indicates whether ransomware detection is based on new file types not seen before at the volume level. This detection method is based only on the file extension not on the file entropy. Some variants of ransomware modify the data such that the file entropy remains unchanged. This method helps in detecting those ransomware but there is a possibility of false positives. Note that ransomware detection is also done based on combined file extension and file entropy and this method of detection is always enabled and has no dependency on this parameter.

`[-based-on-file-create-rate {true|false}]` - Is Based on File Create Operation Rate

This parameter indicates whether ransomware detection is based on the file create rate at the volume level. If this is true and the number of files created per timeslot surges by `-file-create-rate-surge -notify-percentage` percentage compared to the historically observed value, then it is considered an attack.

[-based-on-file-rename-rate {true|false}] - Is Based on File Rename Operation Rate

This parameter indicates whether ransomware detection is based on the file rename rate at the volume level. If this is true and the number of files renamed per timeslot surges by `-file-rename-rate-surge-notify-percentage` percentage compared to the historically observed value, then it is considered an attack.

[-based-on-file-delete-rate {true|false}] - Is Based on File Delete Operation Rate

This parameter indicates whether ransomware detection is based on the file delete rate at the volume level. If this is true and the number of files deleted per timeslot surges by `-file-delete-rate-surge-notify-percentage` percentage compared to the historically observed value, then it is considered an attack.

[-relaxing-popular-file-extensions {true|false}] - Is Relaxing Popular File Extensions

This parameter indicates whether ransomware detection is based on commonly used extensions. If true, then a predetermined commonly used extension, such as .mp3, is considered safe. If false, only those file extensions observed during the dry-run state are considered safe; any extension not observed during the dry-run state but observed later is a suspected ransomware attack, even if it is a commonly used extension.

[-high-entropy-data-surge-notify-percentage <integer>] - High Entropy Data Surge Notify Percentage

This parameter displays the surge value that is considered safe in the overall incoming data at the volume level.

[-file-create-rate-surge-notify-percentage <integer>] - File Create Operation Rate Surge Notify Percentage

This parameter displays the surge rate that is considered safe for file create operations at the volume level.

[-file-delete-rate-surge-notify-percentage <integer>] - File Delete Operation Rate Surge Notify Percentage

This parameter displays the surge rate that is considered safe for file delete operations at the volume level.

[-file-rename-rate-surge-notify-percentage <integer>] - File Rename Operation Rate Surge Notify Percentage

This parameter displays the surge rate that is considered safe for file rename operations at the volume level.

[-never-seen-before-file-extn-count-notify-threshold <integer>] - Never Seen before File Extension Count Notify Threshold

This parameter displays the threshold value of number of files observed with a new file extension not seen before for create/rename operations.

[-never-seen-before-file-extn-duration-in-hours <integer>] - Never Seen before File Extension Duration in Hours

This parameter displays the duration for new file extensions not seen before, in hours. If a new file extension is observed and `-never-seen-before-file-extn-count-notify-threshold` number of files are created/renamed with this new file extension for this duration, then it is reported as an attack.

Examples

The following example displays attack detection parameter information of a volume.

```
cluster1::> security anti-ransomware volume attack-detection-parameters
show -vserver vs1 -volume vol1
                                            Vserver Name : vs1
                                             Volume Name : vol1
            Is Detection Based on High Entropy Data Rate? : true
  Is Detection Based on Never Seen before File Extension? : true
                  Is Detection Based on File Create Rate? : true
                  Is Detection Based on File Rename Rate? : true
                  Is Detection Based on File Delete Rate? : true
           Is Detection Relaxing Popular File Extensions? : true
                High Entropy Data Surge Notify Percentage : 100
                 File Create Rate Surge Notify Percentage : 100
                 File Rename Rate Surge Notify Percentage : 100
                 File Delete Rate Surge Notify Percentage : 100
 Never Seen before File Extensions Count Notify Threshold : 20
       Never Seen before File Extensions Duration in Hour : 24

cluster1::> security anti-ransomware volume attack-detection-parameters
modify -vserver vs1 -volume vol1 -file-delete-rate-surge-notify-percentage
25

cluster1::> security anti-ransomware volume attack-detection-parameters
show -vserver vs1 -volume vol1
                                            Vserver Name : vs1
                                             Volume Name : vol1
            Is Detection Based on High Entropy Data Rate? : true
  Is Detection Based on Never Seen before File Extension? : true
                  Is Detection Based on File Create Rate? : true
                  Is Detection Based on File Rename Rate? : true
                  Is Detection Based on File Delete Rate? : true
           Is Detection Relaxing Popular File Extensions? : true
                High Entropy Data Surge Notify Percentage : 100
                 File Create Rate Surge Notify Percentage : 100
                 File Rename Rate Surge Notify Percentage : 100
                 File Delete Rate Surge Notify Percentage : 25
 Never Seen before File Extensions Count Notify Threshold : 20
       Never Seen before File Extensions Duration in Hour : 24
```

security anti-ransomware volume attack-detection-parameters show

Show anti-ransomware volume attack detection parameters

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

Description

The `security anti-ransomware volume attack-detection-parameters show` command displays attack detection parameter details of an anti-ransomware enabled volume.

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.

-vserver <Vserver Name> - Vserver Name

This parameter specifies the Vserver of the anti-ransomware enabled volume.

-volume <volume name> - Volume Name

This parameter specifies the anti-ransomware enabled volume for which the attack detection parameters need to be displayed.

[-based-on-high-entropy-data-rate {true|false}] - High Entropy Data Rate at Volume Level

This parameter displays whether ransomware detection is based on a high entropy data rate at the volume level. Ransomware detection is also done based on high entropy data rate at the file level and this method of detection is always enabled and has no dependency on this parameter.

[-based-on-never-seen-before-file-extension {true|false}] - Never Seen before File Extension

This parameter indicates whether ransomware detection is based on new file types not seen before at the volume level. This detection method is based only on the file extension not on file entropy. Some variants of ransomware modify the data such that the file entropy remains unchanged. This method helps in detecting those ransomware but there is a possibility of false positives. Note that ransomware detection is also done based on combined file extension and file entropy and this method of detection is always enabled and has no dependency on this parameter.

[-based-on-file-create-rate {true|false}] - Is Based on File Create Operation Rate

This parameter displays whether ransomware detection is based on the file create rate at the volume level. If this is true and the number of files created per timeslot surges by `-file-create-rate-surge -notify-percentage` percentage compared to the historically observed value, then it is considered an attack.

[-based-on-file-rename-rate {true|false}] - Is Based on File Rename Operation Rate

This parameter displays whether ransomware detection is based on the file rename rate at the volume level. If this is true and the number of files renamed per timeslot surges by `-file-rename-rate-surge -notify-percentage` percentage compared to the historically observed value, then it is considered an attack.

[-based-on-file-delete-rate {true|false}] - Is Based on File Delete Operation Rate

This parameter displays whether ransomware detection is based on the file delete rate at the volume level. If this is true and the number of files deleted per timeslot surges by `-file-delete-rate-surge -notify-percentage` percentage compared to the historically observed value, then it is considered an

attack.

[-relaxing-popular-file-extensions {true|false}] - Is Relaxing Popular File Extensions

This parameter displays whether ransomware detection is based on commonly used extensions. If true, then a predetermined commonly used extension, such as .mp3, is considered safe. If false, only those file extensions observed during the dry run state are considered safe; any extension not observed during the dry-run state but observed later is suspected as a ransomware attack, even if it is a commonly used extension.

[-high-entropy-data-surge-notify-percentage <integer>] - High Entropy Data Surge Notify Percentage

This parameter displays the surge value that is considered safe in the overall incoming data at the volume level.

[-file-create-rate-surge-notify-percentage <integer>] - File Create Operation Rate Surge Notify Percentage

This parameter displays the surge rate that is considered safe for file create operations at the volume level.

[-file-delete-rate-surge-notify-percentage <integer>] - File Delete Operation Rate Surge Notify Percentage

This parameter displays the surge rate that is considered safe for file delete operations at the volume level.

[-file-rename-rate-surge-notify-percentage <integer>] - File Rename Operation Rate Surge Notify Percentage

This parameter displays the surge rate that is considered safe for file rename operations at the volume level.

[-never-seen-before-file-extn-count-notify-threshold <integer>] - Never Seen before File Extension Count Notify Threshold

This parameter displays the threshold value of new file extensions not seen before for create/rename operations.

[-never-seen-before-file-extn-duration-in-hours <integer>] - Never Seen before File Extension Duration in Hours

This parameter displays the duration for new file extensions not seen before, in hours. If a new file extension is observed and `-never-seen-before-file-extn-count-notify-threshold` number of files are created/renamed with this new file extension for this duration, then it is reported as an attack.

Examples

The following example displays attack detection parameter information of a volume.

```
cluster1::> security anti-ransomware volume attack-detection-parameters
show -vserver vs1 -volume vol1
                      Vserver Name : vs1
                       Volume Name : vol1
      Is Detection Based on High Entropy Data Rate? : true
  Is Detection Based on Never Seen before File Extension? : true
          Is Detection Based on File Create Rate? : true
          Is Detection Based on File Rename Rate? : true
          Is Detection Based on File Delete Rate? : true
       Is Detection Relaxing Popular File Extensions? : true
        High Entropy Data Surge Notify Percentage : 100
         File Create Rate Surge Notify Percentage : 100
         File Rename Rate Surge Notify Percentage : 100
         File Delete Rate Surge Notify Percentage : 100
 Never Seen before File Extensions Count Notify Threshold : 20
     Never Seen before File Extensions Duration in Hour : 24
```

security anti-ransomware volume auto-switch-to-enable-mode show

Show anti-ransomware volume auto-switch to enable-mode stats

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

Description

The `security anti-ransomware volume auto-switch-to-enable-mode show` command displays stats related to auto-switch to enable-mode of the volumes in dry-run state.

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.

[-vserver <vserver name>] - Vserver Name

This parameter specifies the Vserver of the anti-ransomware enabled volume.

[-volume <volume name>] - Volume Name

This parameter specifies the anti-ransomware enabled volume for which the auto-switch info need to be displayed.

[-total-writes <integer>] - Total writes

Total data writes size so far in dry-run period.

[-dry-run-days-count <integer>] - Dry run days count

Number of days a volume is in learning mode.

[-days-since-new-extn-seen <integer>] - Days since new extn seen

Number of days where no new extension is observed in the volume..

[-dry-run-new-file-count <integer>] - Dry run new file count

Number of files created in a volume during learning period.

[-dry-run-new-file-extension-count <integer>] - Dry run new file extensions count

Number of file extensions observed in a volume during learning period.

Examples

The following example displays attack detection parameter information of a volume.

```
cluster1::> security anti-ransomware volume auto-switch-to-enable-mode
show -vserver vs1 -volume vol1
      Vserver                                                    : vs_1
      Volume                                                     : volt1
      Amount of Write (in KB) Received during Dry-run Mode       :
23920640
      Number of Days Completed in Dry-run Mode                   : 40
      Number of Days Without a New Extension seen in Dry-run Mode  : 2
      Number of Files created during Dry-run Period              : 1991
      Number of File extensions observed in Dry-run period       : 20
```

security anti-ransomware volume event-log modify

Modify anti-ransomware event-log for volumes

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

Description

The `security anti-ransomware volume event-log` command.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver of anti-ransomware enabled volume.

-volume <volume name> - Volume Name

This parameter specifies the anti-ransomware enabled volume.

[-is-enabled-on-new-file-extension-seen {true|false}] - Is Enabled on New File Extension Seen

This parameter indicates whether an alert event should be created when files with new extensions are

created.

[-is-enabled-on-snapshot-copy-creation {true|false}] - Is Enabled on Snapshot Copy Creation

This parameter indicates whether an alert event should be created when snapshots are created.

Examples

```
cluster1::> security anti-ransomware volume event-log modify -vserver vs1
-volume vol1 -is-enabled-on-new-file-extension-seen true -is-enabled-on
-snapshot-copy-creation true

cluster1::> security anti-ransomware volume event-log show -vserver vs1
-volume vol1
Vserver : vs1
                    Volume : vol1
            Is Enabled on
   New File Extension Seen : true
            Is Enabled on
    Snapshot Copy Creation : true
```

security anti-ransomware volume event-log show

Show anti-ransomware event-log for volumes

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

Description

The `security anti-ransomware volume event-log` command.

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.

[-vserver <vserver name>] - Vserver Name

This parameter specifies the Vserver of anti-ransomware enabled volume.

[-volume <volume name>] - Volume Name

This parameter specifies the anti-ransomware enabled volume.

[-is-enabled-on-new-file-extension-seen {true|false}] - Is Enabled on New File Extension Seen

This parameter indicates whether an alert event is generated on creation of files with new extensions.

[-is-enabled-on-snapshot-copy-creation {true|false}] - Is Enabled on Snapshot Copy Creation

This parameter indicates whether an alert event is generated when a snapshot is created.

security anti-ransomware volume space show

Display the details of anti-ransomware space usage

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

Description

This `security anti-ransomware volume space show` displays the space usage by Anti-ransomware feature.

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.

[-vserver <vserver name>] - Vserver Name

This parameter specifies the Vserver on which the volume is located.

[-volume <volume name>] - Volume Name

This parameter specifies the name of the volume whose space usage details are to be shown.

[-space-used-by-snapshot {<integer>[KB|MB|GB|TB|PB]}] - Space Used by snapshots

This parameter shows space usage by Anti-ransomware snapshots.

[-space-used-by-logs {<integer>[KB|MB|GB|TB|PB]}] - Space Used by logs

This parameter shows the space used by the Anti-ransomware logs.

[-total-space-used {<integer>[KB|MB|GB|TB|PB]}] - Total space used by anti-ransomware

This parameter shows the total space used by the Anti-ransomware feature.

[-no-of-snapshot <integer>] - Number of Anti-ransomware Snapshot Copies

This parameter shows the total count of the Anti-ransomware snapshots.

Examples

The following example shows a sample output for this command:

```
clus1::>> security anti-ransomware volume space show
                       Space Used By Space Used By Total Space Snapshot
Vserver    Volume      Snapshot      logs          Used        Copies
---------- ----------- ----------    ----------    ----------
---------------
vs1        vol1             308KB            8B       308.0KB
2
```

security anti-ransomware volume workload-behavior clear-surge

Clear the observed surge values on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security anti-ransomware volume workload-behavior clear-surge` command clears the observed surge values.

Parameters

`-vserver <vserver name>` - Vserver Name

This parameter specifies the Vserver on which the volume is located.

`-volume <volume name>` - Volume Name

This parameter specifies the name of the volume on which anti-ransomware feature is enabled.

Examples

```
cluster1::> security anti-ransomware volume workload-behavior show
-vserver vs1 -volume vol1
                                          Vserver: vs1
                                           Volume: vol1
                        File Extensions Observed: .txt, .exe, .pdf, .img
              Number of File Extensions Observed: 4
Historical Statstics
              High Entropy Data Write Percentage: 50
   High Entropy Data Write Peak Rate (KB/Minute): 50
              File Create Peak Rate (per Minute): 100
              File Delete Peak Rate (per Minute): 100
              File Rename Peak Rate (per Minute): 100
Surge Observed
                                   Surge Timeline: 09/05/2022 14:01:00
              High Entropy Data Write Percentage: 100
   High Entropy Data Write Peak Rate (KB/Minute): 2000
              File Create Peak Rate (per Minute): 80
```

```
              File Delete Peak Rate (per Minute): -
              File Rename Peak Rate (per Minute): 200
                  Newly Observed File Extensions: .dll, .exec, .js
        Number of Newly Observed File Extensions: 10, 4, 22

cluster1::> security anti-ransomware volume workload-behavior clear-surge
-vserver vs1 -volume vol1

cluster1::> security anti-ransomware volume workload-behavior show
-vserver vs1 -volume vol1
                                         Vserver: vs1
                                          Volume: vol1
                        File Extensions Observed: .txt, .exe, .pdf, .img
              Number of File Extensions Observed: 4
Historical Statstics
              High Entropy Data Write Percentage: 50
   High Entropy Data Write Peak Rate (KB/Minute): 50
              File Create Peak Rate (per Minute): 100
              File Delete Peak Rate (per Minute): 100
              File Rename Peak Rate (per Minute): 100
Surge Observed
                                   Surge Timeline: -
              High Entropy Data Write Percentage: -
   High Entropy Data Write Peak Rate (KB/Minute): -
              File Create Peak Rate (per Minute): -
              File Delete Peak Rate (per Minute): -
              File Rename Peak Rate (per Minute): -
                  Newly Observed File Extensions: .dll, .exec, .js
        Number of Newly Observed File Extensions: 10, 4, 22
```

security anti-ransomware volume workload-behavior show

Display information about the volume's workload-behavior learnt by the analytics algorithm

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This `security anti-ransomware volume workload-behavior show` displays the workload characteristics observed during anti-ransomware monitoring.

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.

-vserver <Vserver Name> - Vserver Name

This parameter specifies the Vserver of the anti-ransomware enabled volume.

-volume <volume name> - Volume Name

This parameter specifies the anti-ransomware enabled volume for which the workload behavior details are displayed.

[-file-extensions-included <text>,…] - List of File Extensions Observed

This parameter displays the list of file extensions observed during anti-ransomware monitoring.

[-total-file-extensions-included <integer>] - Number of File Extensions Observed

This parameter displays the number of file extensions observed during anti-ransomware monitoring.

[-high-entropy-data-write-peak-percent <integer>] - High Entropy Data Write Peak Percentage

This parameter displays the peak historical high entropy data write percentage of the incoming data.

[-high-entropy-data-write-peak-rate <integer>] - High Entropy Data Write Peak Rate (KB/minute)

This parameter displays the peak historical high entropy data write rate.

[-file-create-peak-rate <integer>] - File Create Peak Rate per Minute

This parameter displays the peak historical rate of file create operations in the volume.

[-file-rename-peak-rate <integer>] - File Rename Peak Rate per Minute

This parameter displays the peak historical rate of file rename operations in the volume.

[-file-delete-peak-rate <integer>] - File Delete Peak Rate per Minute

This parameter displays the peak historical rate of file delete operations in the volume.

[-surge-timeline <MM/DD/YYYY HH:MM:SS>] - Surge Timeline

This parameter displays the timeline where a surge was observed in the workload characteristics compared to the historically learned characteristics.

[-surge-high-entropy-data-write-peak-percent <integer>] - High Entropy Data Write Percentage During Surge

This parameter displays the peak percentage value of high entropy data write in the incoming data when the surge was observed.

[-surge-high-entropy-data-write-peak-rate <integer>] - High Entropy Data-write Peak Rate Surge (KB/minute)

This parameter displays the peak rate of high entropy data write when the surge was observed.

[-surge-file-create-peak-rate <integer>] - File Create Peak Rate (per Minute) During Surge

This parameter displays the surge in the peak rate of file create operations.

[-surge-file-delete-peak-rate <integer>] - File Delete Peak Rate (per Minute) During Surge

This parameter displays the surge in the peak rate of file delete operations.

[-surge-file-rename-peak-rate <integer>] - File Rename Peak Rate (per Minute) During Surge

This parameter displays the surge in the peak rate of file rename operations.

[-attack-file-extensions-observed <text>,…] - File Extensions Observed During Attack

This parameter displays the list of file types observed during a suspected ransomware attack.

[-attack-file-extensions-observed-counts <integer>,…] - Number of File Extensions Observed During Attack

This parameter displays the count of various file types observed during a suspected ransomware attack.

Examples

The following example shows sample output for this command:

```
cluster1::> security anti-ransomware volume workload-behavior show
-vserver vs1 -volume vol1
                                        Vserver: vs1
                                         Volume: vol1
                        File Extensions Observed: .ext1, .ext2, .ext3
              Number of File Extensions Observed: 3
Historical Statstics
              High Entropy Data Write Percentage: 50
   High Entropy Data Write Peak Rate (KB/Minute): 50
              File Create Peak Rate (per Minute): 100
              File Delete Peak Rate (per Minute): 100
              File Rename Peak Rate (per Minute): 100
Surge Observed
                                  Surge Timeline: 1/1/2022 01:01:01
              High Entropy Data Write Percentage: 200
   High Entropy Data Write Peak Rate (KB/Minute): 200
              File Create Peak Rate (per Minute): 200
              File Delete Peak Rate (per Minute): 200
              File Rename Peak Rate (per Minute): 200
                  Newly Observed File Extensions: .uk1,.uk2,.uk3
        Number of Newly Observed File Extensions: 1, 2, 3
```

security anti-ransomware volume workload-behavior update-baseline-from-surge

Set the observed surge values as the new baseline on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security anti-ransomware volume workload-behavior update-baseline-from-surge` command sets the observed surge value as new baseline.

Parameters

`-vserver <vserver name>` - Vserver Name

This parameter specifies the Vserver on which the volume is located.

`-volume <volume name>` - Volume Name

This parameter specifies the name of the volume on which anti-ransomware feature is enabled.

Examples

```
cluster1::> security anti-ransomware volume workload-behavior show
-vserver vs1 -volume vol1
                                        Vserver: vs1
                                         Volume: vol1
                       File Extensions Observed: .txt, .exe, .pdf, .img
              Number of File Extensions Observed: 4
Historical Statstics
              High Entropy Data Write Percentage: 50
   High Entropy Data Write Peak Rate (KB/Minute): 50
              File Create Peak Rate (per Minute): 100
              File Delete Peak Rate (per Minute): 100
              File Rename Peak Rate (per Minute): 100
Surge Observed
                                  Surge Timeline: 10/3/2021 14:01:00
              High Entropy Data Write Percentage: 100
   High Entropy Data Write Peak Rate (KB/Minute): 2000
              File Create Peak Rate (per Minute): 80
              File Delete Peak Rate (per Minute): -
              File Rename Peak Rate (per Minute): 200
                  Newly Observed File Extensions: .dll, .exec, .js
        Number of Newly Observed File Extensions: 10, 4, 22

cluster1::> security anti-ransomware volume workload-behavior update-
baseline-from-surge -vserver vs1 -volume vol1

cluster1::> security anti-ransomware volume workload-behavior show
-vserver vs1 -volume vol1
                                        Vserver: vs1
                                         Volume: vol1
                       File Extensions Observed: .txt, .exe, .pdf, .img
              Number of File Extensions Observed: 4
Historical Statstics
```

```
               High Entropy Data Write Percentage: 100
   High Entropy Data Write Peak Rate (KB/Minute): 2000
               File Create Peak Rate (per Minute): 180
               File Delete Peak Rate (per Minute): 100
               File Rename Peak Rate (per Minute): 200
 Surge Observed
                                   Surge Timeline: -
               High Entropy Data Write Percentage: -
   High Entropy Data Write Peak Rate (KB/Minute): -
               File Create Peak Rate (per Minute): -
               File Delete Peak Rate (per Minute): -
               File Rename Peak Rate (per Minute): -
                   Newly Observed File Extensions: .dll, .exec, .js
        Number of Newly Observed File Extensions: 10, 4, 22
```

security anti-ransomware vserver event-log modify

Modify anti ransomware event log configuration for Vservers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security anti-ransomware vserver event-log modify` command modifies the anti ransomware event log configuration of a specified Vserver

Parameters

`-vserver <vserver name>` - Vserver

Vserver whose anti ransomware event log configuration will be modified.

`[-is-enabled-on-new-file-extension-seen {true|false}]` - Is Enabled on New File Extension Seen

This parameter indicates whether an alert event should be created when files with new extensions are created.

`[-is-enabled-on-snapshot-copy-creation {true|false}]` - Is Enabled on Snapshot Copy Creation

This parameter indicates whether an alert event should be created when snapshots are created.

security anti-ransomware vserver event-log show

Show anti ransomware event log configuration for Vservers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security anti-ransomware vserver event-log show` command displays anti ransomware

event log configuration for a Vserver.

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.

[-vserver <vserver name>] - Vserver

If this parameter is specified, the command displays anti ransomware event log information about the specified Vserver.

[-is-enabled-on-new-file-extension-seen {true|false}] - Is Enabled on New File Extension Seen

If this parameter is specified, the command displays information only about the Vservers that match the specified value for new file extension seen event logging.

[-is-enabled-on-snapshot-copy-creation {true|false}] - Is Enabled on Snapshot Copy Creation

If this parameter is specified, the command displays information only about the Vservers that match the specified value for snapshot creation event logging.

Examples

The following example shows a sample output for this command:

```
cluster1::> security anti-ransomware vserver event-log show
Vserver: vs_1
 Is Enabled on New File Extension Seen: true
  Is Enabled on Snapshot Copy Creation: false
```

security audit commands

security audit modify

Set administrative audit logging settings

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security audit modify` command modifies the following audit-logging settings for the management interface:

- Whether get requests for the CLI are audited

- Whether get requests for the ONTAP API (ONTAPI) are audited

Parameters

`[-cliget {on|off}]` - Enable Auditing of CLI Get Operations

This specifies whether get requests for the CLI are audited. The default setting is *`off`* .

`[-httpget {on|off}]` - Enable Auditing of HTTP Get Operations

This specifies whether get requests for the web (HTTP) interface are audited. The default setting is *`off`* .

`[-ontapiget {on|off}]` - Enable Auditing of Data ONTAP API Get Operations

This specifies whether get requests for the ONTAP API (ONTAPI) interface are audited. The default setting is *`off`* .

Examples

The following example turns off auditing of get requests for the CLI interface:

```
cluster1::> security audit modify -cliget off
```

security audit show

Show administrative audit logging settings

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security audit show` command displays the following audit-logging settings for the management interface:

- Whether get requests for the CLI are audited
- Whether get requests for the web (HTTP) interface are audited
- Whether get requests for the ONTAP API (ONTAPI) are audited

Audit log entries are written to the 'audit' log, viewable via the 'security audit log show' command.

Examples

The following example displays the audit-logging settings for the management interface:

```
cluster1::> security audit show
              Auditing State for
    Operation Get Requests
    --------- ------------------
          CLI off
         HTTP off
       ONTAPI off
```

security audit log show

Display audit entries merged from multiple nodes in the cluster

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security audit log show` command displays cluster-wide audit log messages. Messages from each node are interleaved in chronological order.

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.

| [-detail]

This display option shows the individual fields of the audit record.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-timestamp <Date>] - Log Entry Timestamp

Selects the entries that match the specified input for timestamp. This will be in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec> <year> in the local timezone.

[-node {<nodename>|local}] - Node

Selects the entries that match the specified input for node.

[-entry <text>] - Log Message Entry

Selects the entries that match the specified input for entry.

[-session-id <text>] - Session ID

This is the "session id" for this audit record. Each ssh/console session is assigned a unique session ID. Each ZAPI/HTTP/SNMP request is assigned a uniqueue session ID

[-command-id <text>] - Command ID

This is useful with ssh/console sessions. Each command in a session is assigned a unique command ID. Each ZAPI/HTTP/SNMP request does not have a command ID.

[-application <text>] - Protocol

This is the application used to connect to the cluster. Possible values include the following: internal, console, ssh, http, ontapi, snmp, rsh, telnet, service-processor

[-location <text>] - Remote user location

The remote IP address or remote access point.

[-vserver <text>] - Vserver name

Storage Virtual Machine name

[-username <text>] - Username

Username

[-input <text>] - Command being executed

The operation being attempted

[-state {Pending|Success|Error}] - State of this audit request

State of this request

[-message <text>] - Additional information and/or error message

Additional information which may be error or informative message.

Examples

The following example displays specific fields based on a custom query:

```
cluster1::> security audit log show -fields application, location, state,
input, message -location 10.60.* -state Error|Success -input v*|st*
-timestamp >"Jul 10 12:00:00 2020"
timestamp                  node  application location     input
state   message
-------------------------- ----- ----------- ------------
------------------------------------------ ------- -------
"Fri Jul 17 11:32:44 2020" node1 ssh         10.60.250.79 storage
aggregate create test -diskcount 5 Success -
"Fri Jul 17 11:36:47 2020" node1 ssh         10.60.250.79 vserver create
vs1                         Success -
"Fri Jul 17 11:37:33 2020" node1 ssh         10.60.250.79 volume create
vol1                         Error   One of the following parameters is
required: -aggregate, -aggr-list, -auto-provision-as
"Fri Jul 17 11:38:08 2020" node1 ssh         10.60.250.79 volume create
vol1 -aggregate test         Success -
Some more examples for -timestamp usage:
cluster1::> security audit log show -timestamp "Mon Jan 03 18:37:05 2022"
Time                       Node         Audit Message
------------------------  -----------  -----------------------
Mon Jan 03 18:37:05 2022  node1
```

```
                                        [kern_audit:info:988] mlogd:
started

cluster1::> security audit log show -timestamp Mon Jan 03 *
Time                      Node         Audit Message
------------------------  -----------  -----------------------
Mon Jan 03 18:37:05 2022  node1
                                        [kern_audit:info:988] mlogd:
started
Mon Jan 03 18:37:06 2022  node2
                                        [kern_audit:info:988] mlogd:
started
Mon Jan 03 18:41:25 2022  node1
                                        [kern_audit:info:977] mlogd:
started
Mon Jan 03 18:41:25 2022  node2
                                        [kern_audit:info:977] mlogd:
started

cluster1::> security audit log show -timestamp Mon Jan 03 18:37*
Time                      Node         Audit Message
------------------------  -----------  -----------------------
Mon Jan 03 18:37:05 2022  node1
                                        [kern_audit:info:988] mlogd:
started
Mon Jan 03 18:37:06 2022  node2
                                        [kern_audit:info:988] mlogd:
started
2 entries were displayed.
```

security certificate commands

security certificate azure-install

Install a Digital Certificate from Azure Key Vault

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security certificate azure-install` command downloads and installs digital security certificates signed by a certificate authority (CA) and the public key certificate of the root CA stored on Azure Key Vault (AKV). With FIPS enabled, the following restrictions apply to the certificate getting installed. server/client/server-ca/client-ca: Key size >= 2048,server/client: Hash function (No MD-5, No SHA-1),server-ca/client-ca: (Intermediate CA), Hash Function (No MD-5, No SHA-1), server-ca/client-ca: (Root CA), Hash Function (No MD-5)

Parameters

`-vserver <Vserver Name>` - Name of Vserver

This specifies the Vserver that contains the certificate.

`-cert-name <text>` - Certificate Name

This specifies the system's internal identifier for the certificate. It must be unique within a Vserver. If not provided, it is automatically generated by the system.

`-type <type of certificate>` - Type of Certificate

This specifies the certificate type. Valid values are the following:

- *`server`* - includes server certificates and intermediate certificates.
- *`client-ca`* - includes the public key certificate for the root CA of the SSL client
- *`server-ca`* - includes the public key certificate for the root CA of the SSL server to which ONTAP is a client
- *`client`* - includes a self-signed or CA-signed digital certificate and private key to be used for ONTAP as an SSL client

`-key-vault-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - Deployed Azure Key Vault DNS Name

The DNS name of the deployed AKV.

`-client-id <text>` - Application (Client) ID of Deployed Azure Application

The ID of the client.

`-tenant-id <text>` - Directory (Tenant) ID of Deployed Azure Application

The ID of the tenant.

`-authentication-method <AKV Authentication Method>` - AKV Authentication Method

Use this parameter to specify the authentication method.

`[-oauth-host <text>]` - Open Authorization Host Name

The hostname of the OAuth server.

`[-proxy-type {http|https}]` - Proxy Type

Proxy Type.

`[-proxy-host <text>]` - Proxy Host

Proxy hostname.

`[-proxy-port <integer>]` - Proxy Port

Proxy port.

`[-proxy-username <text>]` - Proxy Username

Proxy username.

[-proxy-password <text>] - Proxy Password

Proxy password.

[-timeout <integer>] - AKV Connection Timeout in Seconds

AKV Connection Timeout in Seconds.

[-verify-host {true|false}] - Verify the identity of the AKV host

Set to true to verify the identity of the AKV host name.

Examples

This example installs a CA-signed certificate (along with intermediate certificates) for a Vserver named vs0.

```
cluster-1::> security certificate azure-install -vserver vs0 -type client
-client-id client1 -tenant-id tenant1 -key-vault-uri
https://samplevault.vault.azure.net -cert-name certname

Enter the {0} for Azure Key Vault:
```

security certificate create

Create and Install a Self-Signed Digital Certificate

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security certificate create` command creates and installs a self-signed digital certificate, which can be used for server authentication, for signing other certificates by acting as a certificate authority (CA), or for ONTAP as an SSL client. The certificate function is selected by the -type field. Self-signed digital certificates are not as secure as certificates signed by a CA. Therefore, they are not recommended in a production environment.

Parameters

-vserver <Vserver Name> - Name of Vserver

This specifies the name of the Vserver on which the certificate will exist.

-common-name <FQDN or Custom Common Name> - FQDN or Custom Common Name

This specifies the desired certificate name as a fully qualified domain name (FQDN) or custom common name or the name of a person. The supported characters, which are a subset of the ASCII character set, are as follows:

- Letters a through z, A through Z
- Numbers 0 through 9
- Asterisk (*), period (.), underscore (_) and hyphen (-)

The common name must not start or end with a "-" or a ".". The maximum length is 253 characters.

`-type <type of certificate>` - Type of Certificate

This specifies the certificate type. Valid values are the following:

- *`server`* - creates and installs a self-signed digital certificate and intermediate certificates to be used for server authentication
- *`root-ca`* - creates and installs a self-signed digital certificate to sign other certificates by acting as a certificate authority (CA)
- *`client`* - includes a self-signed digital certificate and private key to be used for ONTAP as an SSL client

`[-subtype <kmip-cert>]` - (DEPRECATED)-Certificate Subtype

> This parameter has been deprecated in ONTAP 9.6 and may be removed in a future release of ONTAP.

```
This specifies a certificate subtype. This optional parameter can have
an empty value (the default). The only valid value is as follows:
```

- *`kmip-cert`* - this is a Key Management Interoperability Protocol (KMIP) certificate

`[-cert-name <text>]` - Unique Certificate Name

This specifies the system's internal identifier for the certificate. It must be unique within a Vserver. If not provided, it is automatically generated by the system.

`-size <size of requested certificate in bits>` - Size of Requested Certificate in Bits

This specifies the number of bits in the private key. The larger the value, the more secure is the key. The default is 2048. Possible values include *`512`* , *`1024`* , *`1536`* , *`2048`* and *`3072`* when the "FIPS Mode" in "security config" is false. When the "FIPS Mode" is true, the possible values are *`2048`* and *`3072`* .

`-country <text>` - Country Name

This specifies the country where the Vserver resides. The country name is a two-letter code. The default is US. Here is the list of country codes:
Country Codes

`-state <text>` - State or Province Name

This specifies the state or province where the Vserver resides.

`-locality <text>` - Locality Name

This specifies the locality where the Vserver resides. For example, the name of a city.

`-organization <text>` - Organization Name

This specifies the organization where the Vserver resides. For example, the name of a company.

`-unit <text>` - Organization Unit

This specifies the unit where the Vserver resides. For example, the name of a section or a department within a company.

`-email-addr <mail address>` - Contact Administrator's Email Address

This specifies the email address of the contact administrator for the Vserver.

`-expire-days <integer>` - Number of Days until Expiration

This specifies the number of days until the certificate expires. The default value is 365 days. Possible values are between *1* and *3652* .

`-protocol <protocol>` - Protocol

This specifies the protocol type. This parameter currently supports only the SSL protocol type. The default is SSL.

`-hash-function <hashing function>` - Hashing Function

This specifies the cryptographic hashing function for signing the certificate. The default is SHA256. Possible values include *SHA256* , *SHA224* , *SHA384* and *SHA512* .

Examples

This example creates a server type, self-signed digital certificate for a Vserver named vs0 at a company whose custom common name is *www.example.com* and whose Vserver name is vs0.

```
cluster1::> security certificate create -vserver vs0 -common-name
www.example.com -type server
```

This example creates a root-ca type, self-signed digital certificate with a 2048-bit private key generated by the SHA256 hashing function that will expire in 365 days for a Vserver named vs0 for use by the Software group in IT at a company whose custom common name is *www.example.com* , located in Sunnyvale, California, USA. The email address of the contact administrator who manages the Vserver is *web@example.com* .

```
cluster1::> security certificate create -vserver vs0 -common-name
www.example.com -type root-ca -size 2048 -country US -state California
-locality Sunnyvale -organization IT -unit Software -email-addr
web@example.com -expire-days 365 -hash-function SHA256
```

This example creates a client type of self-signed digital certificate for a Vserver named vs0 at a company that uses ONTAP as an SSL client. The company's custom common name is *www.example.com* and its Vserver name is vs0.

```
cluster1::> security certificate create -vserver vs0 -common-name
www.example.com -type client -size 2048 -country US -state California
-locality Sunnyvale -organization IT -unit Software -email-addr
web@example.com -expire-days 365 -hash-function SHA256
```

security certificate delete

Delete an Installed Digital Certificate

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command deletes an installed digital security certificate.

Parameters

`-vserver <Vserver Name>` - Name of Vserver

This specifies the Vserver that contains the certificate.

`-common-name <FQDN or Custom Common Name>` - FQDN or Custom Common Name

This specifies the desired certificate name as a fully qualified domain name (FQDN) or custom common name or the name of a person. The supported characters, which are a subset of the ASCII character set, are as follows:

- Letters a through z, A through Z
- Numbers 0 through 9
- Asterisk (*), period (.), underscore (_) and hyphen (-)

The common name must not start or end with a "-" or a ".". The maximum length is 253 characters.

`[-serial <text>]` - Serial Number of Certificate

This specifies the certificate serial number.

`-ca <text>` - Certificate Authority

This specifies the certificate authority (CA).

`-type <type of certificate>` - Type of Certificate

This specifies the certificate type. Valid values are the following:

- *`server`* - includes server certificates and intermediate certificates
- *`root-ca`* - includes a self-signed digital certificate to sign other certificates by acting as a certificate authority (CA)
- *`client-ca`* - includes the public key certificate for the root CA of the SSL client. If this client-ca certificate is created as part of a root-ca, it will be deleted along with the corresponding deletion of the root-ca.
- *`server-ca`* - includes the public key certificate for the root CA of the SSL server to which ONTAP is a client. If this server-ca certificate is created as part of a root-ca, it will be deleted along with the corresponding deletion of the root-ca.
- *`client`* - includes a public key certificate and private key to be used for ONTAP as an SSL client

`[-subtype <kmip-cert>]` - (DEPRECATED)-Certificate Subtype

This parameter has been deprecated in ONTAP 9.6 and may be removed in a future release of ONTAP.

```
This specifies a certificate subtype. This optional parameter can have
an empty value (the default). The only valid value is as follows:
```

- *kmip-cert* - this is a Key Management Interoperability Protocol (KMIP) certificate

[-cert-name <text>] - Unique Certificate Name

This specifies the system's internal identifier for the certificate. It is unique within a Vserver.

Examples

This example deletes a root-ca type digital certificate for a Vserver named vs0 in a company named *www.example.com* with serial number 4F57D3D1.

```
cluster1::> security certificate delete -vserver vs0 -common-name
www.example.com -ca www.example.com -type root-ca -serial 4F57D3D1
```

security certificate generate-csr

Generate a Digital Certificate Signing Request

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command generates a digital certificate signing request and displays it on the console. A certificate signing request (CSR or certification request) is a message sent to a certificate authority (CA) to apply for a digital identity certificate.

Parameters

[-common-name <text>] - FQDN or Custom Common Name

This specifies the desired certificate name as a fully qualified domain name (FQDN) or custom common name or the name of a person. The supported characters, which are a subset of the ASCII character set, are as follows:

- Letters a through z, A through Z
- Numbers 0 through 9
- Asterisk (*), period (.), underscore (_) and hyphen (-)

The common name must not start or end with a "-" or a ".". The maximum length is 253 characters.

{ [-size <size of requested certificate in bits>] - (DEPRECATED)-Size of Requested Certificate in Bits

This specifies the number of bits in the private key. A larger size value provides for a more secure key. The default is 2048. Possible values include *512* , *1024* , *1536* , and *2048* .

This parameter has been deprecated in ONTAP 9.8 and may be removed in future releases of ONTAP. Use the security-strength parameter instead.

| [-security-strength <bits of security strength>] - Security Strength in Bits }

Use this parameter to specify the minimum security strength of the certificate in bits. The security bits mapping to RSA and ECDSA key length, in bits, are as follows:

```
Size      RSA Key Length       Elliptic Curve Key Length
112       2048                 224
128       3072                 256
192       4096                 384
```

```
Note: FIPS supported values are restricted to 112 and 128. For ECDSA,
TLSv1.3 requires key length of 256 or greater.
```

[-algorithm <Asymmetric key generation algorithm>] - Asymmetric Encryption Algorithm

Use this parameter to specify the asymmetric encryption algorithm to use for generating the public/private key for the certificate signing request. Algorithm values can be RSA or EC. Default value is RSA.

[-country <text>] - Country Name

This specifies the country where the Vserver resides. The country name is a two-letter code. The default is US. Here is the list of country codes:
Country Codes

[-state <text>] - State or Province Name

This specifies the state or province where the Vserver resides.

[-locality <text>] - Locality Name

This specifies the locality where the Vserver resides. For example, the name of a city.

[-organization <text>] - Organization Name

This specifies the organization where the Vserver resides. For example, the name of a company.

[-unit <text>] - Organization Unit

This specifies the unit where the Vserver resides. For example, the name of a section or a department within a company.

[-email-addr <mail address>] - Contact Administrator's Email Address

This specifies the email address of the contact administrator for the Vserver.

[-hash-function <hashing function>] - Hashing Function

This specifies the cryptographic hashing function for signing the certificate. The default is SHA256. Possible values include *SHA224* , *SHA256* , *SHA384* , and *SHA512* .

[-key-usage <Certificate key usage extension>,…] - Key Usage Extension

Use this parameter to specify the key usage extension values. The default values are: *digitalSignature* , *keyEncipherment* . Possible values include:

- digitalSignature
- nonRepudiation
- keyEncipherment
- dataEncipherment
- keyAgreement
- keyCertSigning
- cRLSigning
- encipherOnly
- decipherOnly
- critical

[-extended-key-usage <Certificate extKeyUsage extension>,…] - Extended Key Usage Extension

Use this parameter to specify the extended key usage extension values. The default values are: *serverAuth* , *clientAuth* . Possible values include:

- serverAuth
- clientAuth
- codeSigning
- emailProtection
- timeStamping
- OCSPSigning
- critical

[-rfc822-name <mail address>,…] - Email Address SAN

Use this parameter to specify the Subject Alternate Name extension - a list of rfc822-names (email addresses).

[-uri <text>,…] - URI SAN

Use this parameter to specify the Subject Alternate Name extension - a list of URIs.

[-dns-name <text>,…] - DNS Name SAN

Use this parameter to specify the Subject Alternate Name extension - a list of DNS names.

[-ipaddr <IP Address>,…] - IP Address SAN

Use this parameter to specify the Subject Alternate Name extension - a list of IP addresses.

Examples

This example creates a certificate-signing request with a 2048-bit RSA private key generated by the SHA256 hashing function for use by the Engineering group in IT at a company whose custom common name is

www.example.com , located in Durham, NC, USA. The email address of the contact administrator who manages the Vserver is *web@example.com* The request also specifies the subject alternative names, key-usage and extended-key-usage extensions.

```
cluster-1::> security certificate generate-csr -common-name
www.example.com -algorithm RSA -hash-function SHA256 -security-strength
128 -key-usage critical,digitalSignature,keyEncipherment -extended-key
-usage serverAuth,clientAuth -country US -state NC -locality Durham
-organization IT -unit Engineering -email-addr web@example.com -rfc822
-name example@example.com -dns-name shop.example.com , store.example.com

Certificate Signing Request :
-----BEGIN CERTIFICATE REQUEST-----
MIIEWDCCAsACAQAwgYgxGDAWBgNVBAMTD3d3dy5leGFtcGxlLmNvbTELMAkGA1UE
BhMCVVMxCzAJBgNVBAgTAk5DMQ8wDQYDVQQHEwZEdXJoYW0xCzAJBgNVBAoTAklU
MRQwEgYDVQQLEwtFbmdpbmVlcmluZzEeMBwGCSqGSIb3DQEJARYPd2ViQGV4YW1w
bGUuY29tMIIBojANBgkqhkiG9w0BAQEFAAOCAY8AMIIBigKCAYEAuo86Jg/szhws
ykYiEXvRaf/j2jJArJMoZby9Z/yINsowe30Xbn5wnfvwiwICUCPwD1e3jhK3TrWH
rNRn/+MqE+jQA7yAdufYxD537cDcT46ihkajISe0Ei93yf6IKmvUAvmJvQ3R7Z4E
QCOWHj56yQ+LXj36bYdwa74S8u8lpCs3Ywx8fgrh/v6H0rnlKDQSQuFR35u7ZZym
tRA7EJMY62f9ALgcFNhQPuP6pjc8aP7Tv7BKXAninryDDcoMdW8UczfTPgzCDh5z
S++eNP3s/7cGfRSQ8aXnDTVQLYpusrdDgVwZXXgu+ZPoZuCf2AYBT+/rdq3VkgWu
QM+mGRMB53O0ff4QOi+SVcXSWXq32wzciv1KsW/iB9h2T+kVd/8Z7ESeYLqFxhY+
0nwacskMRGxOuTLgx+XH+/EntjrI4rjF9/ShYCIcy8vqp1OxFaPClu96ebnbiEOu
y6RvCJ2egcM6OeRbHWB5fIJ0ZZ3crdjz/d1z4ktBuG7E4cUYkEvvAgMBAAGggYkw
gYYGCSqGSIb3DQEJDjF5MHcwRgYDVR0RAQH/BDwwOoETZXhhbXBsZUBleGFtcGxl
LmNvbYIQc2hvcC5leGFtcGxlLmNvbYIRc3RvcmUuZXhhbXBsZS5jb20wDgYDVR0P
AQH/BAQDAgWgMB0GA1UdJQQWMBQGCCsGAQUFBwMCBggrBgEFBQcDATANBgkqhkiG
9w0BAQsFAAOCAYEAh0kOsRy5cCTnFRIWBhBrFFvQhpZIlsoeelNW6JlkE0/ULcAj
JevBx8UibY48D2Wn0nEGle9T3ZeDlg+n66xr/OUfsrENm5ORy5Ndvubkkz0t4KF5
Z2SnwPVIcX2b6ID2xhFAny2S58Adwo7uTpLytidqFj026/KcuyVZUEF9HuJcQGE8
+LMfliCkm6rI2h1ncy2sV6vtDo9GlVscTYLghisHp1aTXVPrr6Q+1OM8lTot8i71
DmZ7kRyxCDlu20XxxV+p2cm4QQVHXbw0XrKAOL2jCBBiYOSWM/BvwWIliVGD6NLg
WK7ZpyHSFjDH0pUlqJCIs079W6JDhiYvtB2xizqmg8oyABUESMUckHGeymr92mcO
JbSyeTE66Pek+Gwia6ZMG7jcznfSr31+7dShLix9kjGsKUffHTiZVySaYjny/+Aq
Seg3Fpusq25ki9D/NMnbifXraL+LbX/WNLS3nA79rp3+VcOoGBponT4i1fsxn+Bv
5RTT3nhT8BlcTe1d
-----END CERTIFICATE REQUEST-----

Private Key :
-----BEGIN PRIVATE KEY-----
MIIG/AIBADANBgkqhkiG9w0BAQEFAASCBuYwggbiAgEAAoIBgQC6jzomD+zOHCzK
RiIRe9Fp/+PaMkCskyhlvL1n/Ig2yjB7fRdufnCd+/CLAgJQI/APV7eOErdOtYes
1Gf/4yoT6NADvIB259jEPnftwNxPjqKGRqMhJ7QSL3fJ/ogqa9QC+Ym9DdHtngRA
I5YePnrJD4tePfpth3BrvhLy7yWkKzdjDHx+CuH+/ofSueUoNBJC4VHfm7tlnKa1
EDsQkxjrZ/0AuBwU2FA+4/qmNzxo/tO/sEpcCeKevIMNygx1bxRzN9M+DMIOHnNL
```

```
7540/ez/twZ9FJDxpecNNVAtim6yt0OBXBldeC75k+hm4J/YBgFP7+t2rdWSBa5A
z6YZEwHnc7R9/hA6L5JVxdJZerfbDNyK/Uqxb+IH2HZP6RV3/xnsRJ5guoXGFj7S
fBpyyQxEbE65MuDH5cf78Se2OsjiuMX39KFgIhzLy+qnU7EVo8KW73p5uduIQ67L
pG8InZ6Bwzo55FsdYHl8gnRlndyt2PP93XPiS0G4bsThxRiQS+8CAwEAAQKCAYBW
fqtWFFIVaWi2y3dmJcL840AP3PaxTHURXkVund3FkU6TIncnqoWqKbHnsSHDaDYX
1vJqc3D7lBx4W+5v7DGJE4rGALKK7olIyzGtUJqUZCwkF0Hw0EijmdBvHYyiJmYg
jvN2bJ7lDTspRZaHJS6mY4eZRSEDgST1PyXn7krEZ6kBSju58G/BWt88KyX80s+Y
pIDiLIDg5pVAI2tPDvQhyI+7sqCKZZQm5GpEgB2JDIS+PgzryUWBlSMp1ICcPcgx
rarFZQi1Ne7qrp6FfKvPAO5XLyI0xhgm8fCMJUpxmEb80XY4FeRDzB42a0Z/YL0P
HhpWAI4ZRsDyDd5S7jwLZQ3Hl9WsKvj2/FRU6hWTP+maH/Vel35iLkygfZWUAjNY
F6B0SoBBd9bVeKDODXrD/CwVbuaKZGMaVOenZbczmFUVSi4HZGyqVRxX6WIxVoD0
MZXwWUoWZ32C6II3vp/ReAsouhCnKDKhqfrvH58xF82FTMMXBZ/kDy7k5IySylkC
gcEA4tpiV1eKzC/ft0sPUNmZB/snHfXC+xohzTygCg4L1Rf8zjDnUT/o9D8SRe1/
crkG7ZcjKvIdPz0tatyjyNMsZ9TDISiAJQJ8Et1+jBP0uy2qG+ab+Ub761BR5TX0
O78UcmtEyxaaDZsESWj+qYerG4E7zGZiTscTe2Jma5fPlS1ekyfNzk1GBtya9bIM
r991o/PahSmCz5iPxf4avYM/vQm2p+wIk+o6ZhJIAUlRFrCv8y9lYivQjw+tZA+G
bdE7AoHBANKHg0Jb5BLJmN/5/PLkkELhaZG+UNUngtm46dm/84+sqtdTcUHpqdHv
M/skRYDVERmI50QZ2HmzVC8J+zzs9r01VNNA+Tzcoi3eB3FPdDYPTDtLSzRfsC82
kix8d2uVs+rfmvKwT0XucNvMQjUyYDII7IJln1iIJp2XQZaNleqgyi65kni+6FrQ
EJ9gVD4PtCkX7rKo8csMITe6n+HZIzFpOY6BX0HU/4VGa+RQHGfGIdfKDOJ5AtyG
RPYVvZ1E3QKBwE520sT7FpsBhBPV9no0iWXlTOZj9wj7RO3EJmbT7OvL3DlFWP0V
afHxTtS5DPgVX3wWZqeYDt2sv2TS5CO2Rwmy4bs6Uvh6H4g27GpvDJshdFEqNpDG
KKR/p5PsUYnI0b2xtJ26N5a1I4pwsoTY1CozTQep8h7lZKusoVhdrgMfKjMj9V+C
AtKkw0RwTUsXs4z973tXnFNJpZEKDx21o/oyvebfESh4P7LGZ/lp7o42luU6Y4rN
NNoGxiZx6EFbuQKBwGbMltJTTmXCHKzZQ6NS6gJOUR9CX/QFLAamHUIfUY3JUU59
RyNZNnv1IluyVWHYKFZgnBSLzkF2yFeDtzMDvmObZAUXh9wpG+Prs5SnqGYxSBb3
6Av14XDcY7nnOOTGn6jDcMSqRLsv99nLvlR9ea1U4C+38XvoV3rB/dvG3PpJcxAn
uxbMmWamjEdWYSxAvMcIEZ0Zk5+DF8E/loxQW7fn2pv0HhBmMjLgtRQx7fzaKXJW
Db6UOkp2IbxL11+w3QKBwDloDgwB7ukGyFHf3RKy3YX0en1WGBesXONf1m2fjwOU
nojccfaGwAUdb6m60JuZFhJ3qZ4ecoloY4GxIKV5krvBg1buow/aqDDkKmVVYNO6
FUuXp+BbTBSxjfftSaog7y5Db5aecLXU5FLE+sVlrhp17s9h8Ur+O04SytSVh9JS
SkzHYv+4GybZqmOeF2U+whib8JXD2bJkSfNI1dZZhKVqoTUQfEAE3VFY0EHkVQwk
rLHmjspsUjKc4BKfVRGWJg==
-----END PRIVATE KEY-----

Note: Keep a copy of your certificate request and private key for future
reference.
```


security certificate install

Install a Digital Certificate

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security certificate install` command installs digital security certificates signed by a certificate authority (CA) and the public key certificate of the root CA. Digital security certificates also include the intermediate certificates to construct the chain for server certificates (the server type), client-side root CA certificates (the client-ca type), or server-side root CA certificates (the server-ca type). with FIPS enabled, the following restrictions apply to the certificate getting installed. server/client/server-ca/client-ca: Key size >= 2048,server/client: Hash function (No MD-5, No SHA-1),server-ca/client-ca: (Intermediate CA), Hash Function (No MD-5, No SHA-1), server-ca/client-ca: (Root CA), Hash Function (No MD-5)

Parameters

`-vserver <Vserver Name>` - Name of Vserver

This specifies the Vserver that contains the certificate.

`-type <type of certificate>` - Type of Certificate

This specifies the certificate type. Valid values are the following:

- *`server`* - includes server certificates and intermediate certificates.
- *`client-ca`* - includes the public key certificate for the root CA of the SSL client
- *`server-ca`* - includes the public key certificate for the root CA of the SSL server to which ONTAP is a client
- *`client`* - includes a self-signed or CA-signed digital certificate and private key to be used for ONTAP as an SSL client

`[-subtype <kmip-cert>]` - (DEPRECATED)-Certificate Subtype

> This parameter has been deprecated in ONTAP 9.6 and may be removed in a future release of ONTAP.

```
This specifies a certificate subtype. This optional parameter can have
an empty value (the default). The only valid value is as follows:
```

- *`kmip-cert`* - this is a Key Management Interoperability Protocol (KMIP) certificate

`[-cert-name <text>]` - Unique Certificate Name

This specifies the system's internal identifier for the certificate. It must be unique within a Vserver. If not provided, it is automatically generated by the system.

Examples

This example installs a CA-signed certificate (along with intermediate certificates) for a Vserver named vs0.

```
cluster1::> security certificate install -vserver vs0 -type server
Enter certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIB8TCCAZugAwIBAwIBADANBgkqhkiG9w0BAQQFADBfMRMwEQYDVQQDEwpuZXRh
cHAuY29tMQswCQYDVQQGEwJVUzEJMAcGA1UECBMAMQkwBwYDVQQHEwAxCTAHBgNV
```

```
BAoTADEJMAcGA1UECxMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcNMTAwNDI2MTk0OTI4
WhcNMTAwNTI2MTk0OTI4WjBfMRMwEQYDVQQDEwpuZXRhcHAuY29tMQswCQYDVQQG
EwJVUzEJMAcGA1UECBMAMQkwBwYDVQQHEwAxCTAHBgNVBAoTADEJMAcGA1UECxMA
MQ8wDQYJKoZIhvcNAQkBFgAwXDANBgkqhkiG9w0BAQEFAANLADBIAkEAyXrK2sry
-----END CERTIFICATE-----
Enter private key: Press <Enter> when done
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAMl6ytrK8nQj82UsWeHOeT8gk0BPX+Y5MLycsUdXA7hXhumHNpvF
C61X2G32Sx8VEa1th94tx+vOEzq+UaqHlt0CAwEAAQJBAMZjDWlgmlm3qIr/n8VT
PFnnZnbVcXVM7OtbUsgPKw+QCCh9dF1jmuQKeDr+wUMWknlDeGrfhILpzfJGHrLJ
z7UCIQDr8d3gOG71UyX+BbFmo/N0uAKjS2cvUU+Y8a8pDxGLLwIhANqa99SuSl8U
DiPvdaKTj6+EcGuXfCXz+G0rfgTZK8uzAiEAr1mnrfYC8KwE9k7A0ylRzBLdUwK9
AvuJDn+/z+H1Bd0CIQDD93P/xpaJETNz53Au49VE5Jba/Jugckrbosd/lSd7nQIg
aEMAzt6qHHT4mndi8Bo8sDGedG2SKx6Qbn2IpuNZ7rc=
-----END RSA PRIVATE KEY-----

Do you want to continue entering root and/or intermediate certificates
{y|n}: y

Enter intermediate certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIE+zCCBGSgAwIBAgICAQ0wDQYJKoZIhvcNAQEFBQAwgbsxJDAiBgNVBAcTG1Zh
bGlDZXJ0IFZhbGlkYXRpb24gTmV0d29yazEXMBUGA1UEChMOVmFsaUNlcnQsIElu
Yy4xNTAzBgNVBAsTLFZhbGlDZXJ0IENsYXNzIDIgUG9saWN5IFZhbGlkYXRpb24g
QXV0aG9yaXR5MSEwHwYDVQQDExhodHRwOi8vd3d3LnZhbGljZXJ0LmNvbS8xIDAe
BgkqhkiG9w0BCQEWEWluZm9AdmFsaWNlcnQuY29tMB4XDTA0MDYyOTE3MDYyMFoX
DTI0MDYyOTE3MDYyMFowYzELMAkGA1UEBhMCVVMxITAfBgNVBAoTGFRoZSBHbyBE
YWRkeSBHcm91cCwgSW5jLjExMC8GA1UECxMoR28gRGFkZHkgQ2xhc3MgMiBDZXJ0
-----END CERTIFICATE-----
Do you want to continue entering root and/or intermediate certificates
{y|n}: y

Enter intermediate certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIC5zCCAlACAQEwDQYJKoZIhvcNAQEFBQAwgbsxJDAiBgNVBAcTG1ZhbGlDZXJ0
IFZhbGlkYXRpb24gTmV0d29yazEXMBUGA1UEChMOVmFsaUNlcnQsIEluYy4xNTAz
BgNVBAsTLFZhbGlDZXJ0IENsYXNzIDIgUG9saWN5IFZhbGlkYXRpb24gQXV0aG9y
aXR5MSEwHwYDVQQDExhodHRwOi8vd3d3LnZhbGljZXJ0LmNvbS8xIDAeBgkqhkiG
9w0BCQEWEWluZm9AdmFsaWNlcnQuY29tMB4XDTk5MDYyNjAwMTk1NFoXDTE5MDYy
NjAwMTk1NFowgbsxJDAiBgNVBAcTG1ZhbGlDZXJ0IFZhbGlkYXRpb24gTmV0d29y
azEXMBUGA1UEChMOVmFsaUNlcnQsIEluYy4xNTAzBgNVBAsTLFZhbGlDZXJ0IENs
YXNzIDIgUG9saWN5IFZhbGlkYXRpb24gQXV0aG9yaXR5MSEwHwYDVQQDExhodHRw
-----END CERTIFICATE-----
Do you want to continue entering root and/or intermediate certificates
{y|n}: n
```

```
You should keep a copy of the private key and the CA-signed digital
certificate
for future reference.
```

This example installs a CA certificate for client authentication for a Vserver named vs0.

```
cluster1::> security certificate install -vserver vs0 -type client-ca

Enter certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIDNjCCAp+gAwIBAgIQNhIilsXjOKUgodJfTNcJVDANBgkqhkiG9w0BAQUFADCB
zjELMAkGA1UEBhMCWkExFTATBgNVBAgTDFdlc3Rlcm4gQ2FwZTESMBAGA1UEBxMJ
Q2FwZSBUb3duMR0wGwYDVQQKExRUaGF3dGUgQ29uc3VsdGluZyBjYzEoMCYGA1UE
CxMfQ2VydGlmaWNhdGlvbiBTZXJ2aWNlcyBEaXZpc2lvbjEhMB8GA1UEAxMYVGhh
d3RlIFByZW1pdW0gU2VydmVyIENBMSgwJgYJKoZIhvcNAQkBFhlwcmVtaXVtLXNl
cnZlckB0aGF3dGUuY29tMB4XDTk2MDgwMTAwMDAwMFoXDTIxMDEwMTIzNTk1OVow
gc4xCzAJBgNVBAYTAlpBMRUwEwYDVQQIEwxXZXN0ZXJuIENhcGUxEjAQBgNVBAcT
-----END CERTIFICATE-----
You should keep a copy of the CA-signed digital certificate for future
reference.
```

This example installs a CA certificate for server authentication for a Vserver named vs0. In this case, ONTAP acts as an SSL client.

```
cluster1::> security certificate install -vserver vs0 -type server-ca

Enter certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIDNjCCAp+gAwIBAgIQNhIilsXjOKUgodJfTNcJVDANBgkqhkiG9w0BAQUFADCB
zjELMAkGA1UEBhMCWkExFTATBgNVBAgTDFdlc3Rlcm4gQ2FwZTESMBAGA1UEBxMJ
Q2FwZSBUb3duMR0wGwYDVQQKExRUaGF3dGUgQ29uc3VsdGluZyBjYzEoMCYGA1UE
CxMfQ2VydGlmaWNhdGlvbiBTZXJ2aWNlcyBEaXZpc2lvbjEhMB8GA1UEAxMYVGhh
d3RlIFByZW1pdW0gU2VydmVyIENBMSgwJgYJKoZIhvcNAQkBFhlwcmVtaXVtLXNl
cnZlckB0aGF3dGUuY29tMB4XDTk2MDgwMTAwMDAwMFoXDTIxMDEwMTIzNTk1OVow
gc4xCzAJBgNVBAYTAlpBMRUwEwYDVQQIEwxXZXN0ZXJuIENhcGUxEjAQBgNVBAcT
-----END CERTIFICATE-----
You should keep a copy of the CA-signed digital certificate for future
reference.
```

security certificate print

Display the contents of a certificate

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the details of either an installed certificate or by reading a certificate from user input.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Use this parameter to specify the Vserver that has the certificate installed.

{ `[-cert-name <text>]` - Installed Certificate Name

Use this parameter to specify the unique name of the installed certificate to read and display.

| `[-cert-uuid <UUID>]` - Installed Certificate UUID }

Use this parameter to specify the unique UUID of the installed certificate to read and display. With no name or UUID specified, the certificate will read and display from user input.

Examples

The following example reads and prints the details of the certificate.

```
cluster1::> security certificate print -vserver vs0 -cert-name
AAACertificateServices
Certificate details:
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 6271844772424770508 (0x570a119742c4e3cc)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=IT, L=Milan, O=Actalis S.p.A./03358520967, CN=Actalis
Authentication Root CA
        Validity
            Not Before: Sep 22 11:22:02 2011 GMT
            Not After : Sep 22 11:22:02 2030 GMT
        Subject: C=IT, L=Milan, O=Actalis S.p.A./03358520967, CN=Actalis
Authentication Root CA
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (4096 bit)
                Modulus:
                    00:a7:c6:c4:a5:29:a4:2c:ef:e5:18:c5:b0:50:a3:
                    6f:51:3b:9f:0a:5a:c9:c2:48:38:0a:c2:1c:a0:18:
                    7f:91:b5:87:b9:40:3f:dd:1d:68:1f:08:83:d5:2d:
                    1e:88:a0:f8:8f:56:8f:6d:99:02:92:90:16:d5:5f:
                    08:6c:89:d7:e1:ac:bc:20:c2:b1:e0:83:51:8a:69:
                    4d:00:96:5a:6f:2f:c0:44:7e:a3:0e:e4:91:cd:58:
                    ee:dc:fb:c7:1e:45:47:dd:27:b9:08:01:9f:a6:21:
                    1d:f5:41:2d:2f:4c:fd:28:ad:e0:8a:ad:22:b4:56:
                    65:8e:86:54:8f:93:43:29:de:39:46:78:a3:30:23:
```

```
                    ba:cd:f0:7d:13:57:c0:5d:d2:83:6b:48:4c:c4:ab:
                    9f:80:5a:5b:3a:bd:c9:a7:22:3f:80:27:33:5b:0e:
                    b7:8a:0c:5d:07:37:08:cb:6c:d2:7a:47:22:44:35:
                    c5:cc:cc:2e:8e:dd:2a:ed:b7:7d:66:0d:5f:61:51:
                    22:55:1b:e3:46:e3:e3:3d:d0:35:62:9a:db:af:14:
                    c8:5b:a1:cc:89:1b:e1:30:26:fc:a0:9b:1f:81:a7:
                    47:1f:04:eb:a3:39:92:06:9f:99:d3:bf:d3:ea:4f:
                    50:9c:19:fe:96:87:1e:3c:65:f6:a3:18:24:83:86:
                    10:e7:54:3e:a8:3a:76:24:4f:81:21:c5:e3:0f:02:
                    f8:93:94:47:20:bb:fe:d4:0e:d3:68:b9:dd:c4:7a:
                    84:82:e3:53:54:79:dd:db:9c:d2:f2:07:9b:2e:b6:
                    bc:3e:ed:85:6d:ef:25:11:f2:97:1a:42:61:f7:4a:
                    97:e8:8b:b1:10:07:fa:65:81:b2:a2:39:cf:f7:3c:
                    ff:18:fb:c6:f1:5a:8b:59:e2:02:ac:7b:92:d0:4e:
                    14:4f:59:45:f6:0c:5e:28:5f:b0:e8:3f:45:cf:cf:
                    af:9b:6f:fb:84:d3:77:5a:95:6f:ac:94:84:9e:ee:
                    bc:c0:4a:8f:4a:93:f8:44:21:e2:31:45:61:50:4e:
                    10:d8:e3:35:7c:4c:19:b4:de:05:bf:a3:06:9f:c8:
                    b5:cd:e4:1f:d7:17:06:0d:7a:95:74:55:0d:68:1a:
                    fc:10:1b:62:64:9d:6d:e0:95:a0:c3:94:07:57:0d:
                    14:e6:bd:05:fb:b8:9f:e6:df:8b:e2:c6:e7:7e:96:
                    f6:53:c5:80:34:50:28:58:f0:12:50:71:17:30:ba:
                    e6:78:63:bc:f4:b2:ad:9b:2b:b2:fe:e1:39:8c:5e:
                    ba:0b:20:94:de:7b:83:b8:ff:e3:56:8d:b7:11:e9:
                    3b:8c:f2:b1:c1:5d:9d:a4:0b:4c:2b:d9:b2:18:f5:
                    b5:9f:4b
                Exponent: 65537 (0x10001)
        X509v3 extensions:
            X509v3 Subject Key Identifier:

52:D8:88:3A:C8:9F:78:66:ED:89:F3:7B:38:70:94:C9:02:02:36:D0
            X509v3 Basic Constraints: critical
                CA:TRUE
            X509v3 Authority Key Identifier:

keyid:52:D8:88:3A:C8:9F:78:66:ED:89:F3:7B:38:70:94:C9:02:02:36:D0
X509v3 Key Usage: critical
                Certificate Sign, CRL Sign
    Signature Algorithm: sha256WithRSAEncryption
         0b:7b:72:87:c0:60:a6:49:4c:88:58:e6:1d:88:f7:14:64:48:
         a6:d8:58:0a:0e:4f:13:35:df:35:1d:d4:ed:06:31:c8:81:3e:
         6a:d5:dd:3b:1a:32:ee:90:3d:11:d2:2e:f4:8e:c3:63:2e:23:
         66:b0:67:be:6f:b6:c0:13:39:60:aa:a2:34:25:93:75:52:de:
         a7:9d:ad:0e:87:89:52:71:6a:16:3c:19:1d:83:f8:9a:29:65:
         be:f4:3f:9a:d9:f0:f3:5a:87:21:71:80:4d:cb:e0:38:9b:3f:
         bb:fa:e0:30:4d:cf:86:d3:65:10:19:18:d1:97:02:b1:2b:72:
```

```
        42:68:ac:a0:bd:4e:5a:da:18:bf:6b:98:81:d0:fd:9a:be:5e:
        15:48:cd:11:15:b9:c0:29:5c:b4:e8:88:f7:3e:36:ae:b7:62:
        fd:1e:62:de:70:78:10:1c:48:5b:da:bc:a4:38:ba:67:ed:55:
        3e:5e:57:df:d4:03:40:4c:81:a4:d2:4f:63:a7:09:42:09:14:
        fc:00:a9:c2:80:73:4f:2e:c0:40:d9:11:7b:48:ea:7a:02:c0:
        d3:eb:28:01:26:58:74:c1:c0:73:22:6d:93:95:fd:39:7d:bb:
        2a:e3:f6:82:e3:2c:97:5f:4e:1f:91:94:fa:fe:2c:a3:d8:76:
        1a:b8:4d:b2:38:4f:9b:fa:1d:48:60:79:26:e2:f3:fd:a9:d0:
        9a:e8:70:8f:49:7a:d6:e5:bd:0a:0e:db:2d:f3:8d:bf:eb:e3:
        a4:7d:cb:c7:95:71:e8:da:a3:7c:c5:c2:f8:74:92:04:1b:86:
        ac:a4:22:53:40:b6:ac:fe:4c:76:cf:fb:94:32:c0:35:9f:76:
        3f:6e:e5:90:6e:a0:a6:26:a2:b8:2c:be:d1:2b:85:fd:a7:68:
        c8:ba:01:2b:b1:6c:74:1d:b8:73:95:e7:ee:b7:c7:25:f0:00:
        4c:00:b2:7e:b6:0b:8b:1c:f3:c0:50:9e:25:b9:e0:08:de:36:
        66:ff:37:a5:d1:bb:54:64:2c:c9:27:b5:4b:92:7e:65:ff:d3:
        2d:e1:b9:4e:bc:7f:a4:41:21:90:41:77:a6:39:1f:ea:9e:e3:
        9f:d0:66:6f:05:ec:aa:76:7e:bf:6b:16:a0:eb:b5:c7:fc:92:
        54:2f:2b:11:27:25:37:78:4c:51:6a:b0:f3:cc:58:5d:14:f1:
        6a:48:15:ff:c2:07:b6:b1:8d:0f:8e:5c:50:46:b3:3d:bf:01:
        98:4f:b2:59:54:47:3e:34:7b:78:6d:56:93:2e:73:ea:66:28:
        78:cd:1d:14:bf:a0:8f:2f:2e:b8:2e:8e:f2:14:8a:cc:e9:b5:
        7c:fb:6c:9d:0c:a5:e1:96
```

security certificate rename

Rename a certificate

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command allows the user to modify the name of an installed digital certificate. This command does not alter the certificate itself.

Parameters

`-vserver <Vserver Name>` - Vserver Name

This specifies the name of the Vserver on which the certificate exists.

`-cert-name <text>` - Existing Certificate Name

This specifies the current name of the certificate.

`-new-name <text>` - New Certificate Name

This specifies the desired name of the certificate. It must be unique among certificates in the Vserver.

Examples

```
cluster1::> security certificate rename -vserver vs0 -cert-name
AAACertificateServices -new-nameAAACertServ
```

security certificate show-generated

Display ONTAP generated certificates

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about the ONTAP generated digital digital certificates. Some details are displayed only when you use the command with the -instance parameter.

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.

[-vserver <Vserver Name>] - Name of Vserver

Selects the Vserver whose digital certificates you want to display.

[-common-name <FQDN or Custom Common Name>] - FQDN or Custom Common Name

Selects the certificates that match this parameter value.

[-serial <text>] - Serial Number of Certificate

Selects the certificates that match this parameter value.

[-ca <text>] - Certificate Authority

Selects the certificates that match this parameter value.

[-type <type of certificate>] - Type of Certificate

Selects the certificates that match this parameter value.

[-subtype <kmip-cert>] - (DEPRECATED)-Certificate Subtype

This parameter has been deprecated in ONTAP 9.6 and may be removed in a future release of ONTAP.

```
Selects the certificate subtype that matches the specified value. The
valid values are as follows:
```

- *`kmip-cert`* - this is a Key Management Interoperability Protocol (KMIP) certificate

[-cert-name <text>] - Unique Certificate Name

This specifies the system's internal identifier for the certificate. It is unique within a Vserver.

[-size <size of requested certificate in bits>] - Size of Requested Certificate in Bits

Selects the certificates that match this parameter value.

[-start <Date>] - Certificate Start Date

Selects the certificates that match this parameter value.

[-expiration <Date>] - Certificate Expiration Date

Selects the certificates that match this parameter value.

[-public-cert <certificate>] - Public Key Certificate

Selects the certificates that match this parameter value.

[-country <text>] - Country Name

Selects the certificates that match this parameter value.

[-state <text>] - State or Province Name

Selects the certificates that match this parameter value.

[-locality <text>] - Locality Name

Selects the certificates that match this parameter value.

[-organization <text>] - Organization Name

Selects the certificates that match this parameter value.

[-unit <text>] - Organization Unit

Selects the certificates that match this parameter value.

[-email-addr <mail address>] - Contact Administrator's Email Address

Selects the certificates that match this parameter value.

[-protocol <protocol>] - Protocol

Selects the certificates that match this parameter value.

[-hash-function <hashing function>] - Hashing Function

Selects the certificates that match this parameter value.

[-self-signed {true|false}] - Self-Signed Certificate

Selects the certificates that match this parameter value.

[-is-root {true|false}] - Is Root CA Certificate?

Selects the certificates that match this parameter value.

[-authority-key-identifier <text>] - Authority Key Identifier

Selects the certificates that match this parameter value.

[-subject-key-identifier <text>] - Subject Key Identifier

Selects the certificates that match this parameter value.

[-rfc822-name <mail address>,…] - Email Address SAN

Selects the certificates that match this parameter value.

[-uri <text>,…] - URI SAN

Selects the certificates that match this parameter value.

[-dns-name <text>,…] - DNS Name SAN

Selects the certificates that match this parameter value.

[-ipaddr <IP Address>,…] - IP Address SAN

Selects the certificates that match this parameter value.

Examples

The examples below display information about ONTAP generated digital certificates.

```
cluster1::> security certificate show-generated

Vserver    Serial Number   Certificate Name                          Type
---------- --------------- ------------------------------------------
---------
vs0        4F4E4D7B         www.example.com         server
    Certificate Authority:  www.example.com
          Expiration Date: Thu Feb 28 16:08:28 2013
```

```
cluster1::> security certificate show-generated -instance
                             Vserver: vs0
                    Certificate Name:  www.example.com
          FQDN or Custom Common Name:  www.example.com
        Serial Number of Certificate: 4F4E4D7B
               Certificate Authority:  www.example.com
                 Type of Certificate: server
 Size of Requested Certificate(bits): 2048
              Certificate Start Date: Fri Apr 30 14:14:46 2010
         Certificate Expiration Date: Sat Apr 30 14:14:46 2011
              Public Key Certificate: -----BEGIN CERTIFICATE-----

MIIDfTCCAmWgAwIBAwIBADANBgkqhkiG9w0BAQsFADBgMRQwEgYDVQQDEwtsYWIu

YWJjLmNvbTELMAkGA1UEBhMCVVMxCTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYD

VQQKEwAxCTAHBgNVBAsTADEPMA0GCSqGSIb3DQEJARYAMB4XDTEwMDQzMDE4MTQ0

BgNVHQ8BAf8EBAMCAQYwHQYDVR0OBBYEFCVG7dYGe51akE14ecaCdL+LOAxUMA0G

CSqGSIb3DQEBCwUAA4IBAQBJlE51pkDY3ZpsSrQeMOoWLteIR+1H0wKZOM1Bhy6Q

+gsE3XEtnN07AE4npjIT0eVP0nI9QIJAbP0uPKaCGAVBSBMoM2mOwbfswI7aJoEh

+XuEoNr0GOz+mltnfhgvl1fT6Ms+xzd3LGZYQTworus2
                                        -----END CERTIFICATE-----
        Country Name (2 letter code): US
  State or Province Name (full name): California
           Locality Name (e.g. city): Sunnyvale
    Organization Name (e.g. company): example
    Organization Unit (e.g. section): IT
        Email Address (Contact Name):  web@example.com
                            Protocol: SSL
                    Hashing Function: SHA256
```

security certificate show-truststore

Display default truststore certificates

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about the default CA certificates that come pre-installed with ONTAP. Some details are displayed only when you use the command with the -instance parameter.

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.

[-vserver <Vserver Name>] - Name of Vserver

Selects the Vserver whose digital certificates you want to display.

[-common-name <FQDN or Custom Common Name>] - FQDN or Custom Common Name

Selects the certificates that match this parameter value.

[-serial <text>] - Serial Number of Certificate

Selects the certificates that match this parameter value.

[-ca <text>] - Certificate Authority

Selects the certificates that match this parameter value.

[-type <type of certificate>] - Type of Certificate

Selects the certificates that match this parameter value.

[-subtype <kmip-cert>] - (DEPRECATED)-Certificate Subtype

This parameter has been deprecated in ONTAP 9.6 and may be removed in a future release of ONTAP.

```
Selects the certificate subtype that matches the specified value. The
valid values are as follows:
```

- *kmip-cert* - this is a Key Management Interoperability Protocol (KMIP) certificate

[-cert-name <text>] - Unique Certificate Name

This specifies the system's internal identifier for the certificate. It is unique within a Vserver.

[-size <size of requested certificate in bits>] - Size of Requested Certificate in Bits

Selects the certificates that match this parameter value.

[-start <Date>] - Certificate Start Date

Selects the certificates that match this parameter value.

[-expiration <Date>] - Certificate Expiration Date

Selects the certificates that match this parameter value.

[-public-cert <certificate>] - Public Key Certificate

Selects the certificates that match this parameter value.

[-country <text>] - Country Name

Selects the certificates that match this parameter value.

[-state <text>] - State or Province Name

Selects the certificates that match this parameter value.

[-locality <text>] - Locality Name

Selects the certificates that match this parameter value.

[-organization <text>] - Organization Name

Selects the certificates that match this parameter value.

[-unit <text>] - Organization Unit

Selects the certificates that match this parameter value.

[-email-addr <mail address>] - Contact Administrator's Email Address

Selects the certificates that match this parameter value.

[-protocol <protocol>] - Protocol

Selects the certificates that match this parameter value.

[-hash-function <hashing function>] - Hashing Function

Selects the certificates that match this parameter value.

[-self-signed {true|false}] - Self-Signed Certificate

Selects the certificates that match this parameter value.

[-is-root {true|false}] - Is Root CA Certificate?

Selects the certificates that match this parameter value.

[-authority-key-identifier <text>] - Authority Key Identifier

Selects the certificates that match this parameter value.

[-subject-key-identifier <text>] - Subject Key Identifier

Selects the certificates that match this parameter value.

[-rfc822-name <mail address>,…] - Email Address SAN

Selects the certificates that match this parameter value.

[-uri <text>,…] - URI SAN

Selects the certificates that match this parameter value.

[-dns-name <text>,…] - DNS Name SAN

Selects the certificates that match this parameter value.

[-ipaddr <IP Address>,…] - IP Address SAN

Selects the certificates that match this parameter value.

Examples

The examples below display information about the pre-installed truststore digital certificates.

```
cluster1::> security certificate show-truststore

Vserver    Serial Number   Certificate Name                          Type
---------- --------------- ------------------------------------------
---------
vs0        4F4E4D7B         www.example.com         server-ca
    Certificate Authority:  www.example.com
          Expiration Date: Thu Feb 28 16:08:28 2013
```

```
cluster1::> security certificate show-truststore -instance
                             Vserver: vs0
                    Certificate Name:  www.example.com
          FQDN or Custom Common Name:  www.example.com
        Serial Number of Certificate: 4F4E4D7B
               Certificate Authority:  www.example.com
                 Type of Certificate: server-ca
 Size of Requested Certificate(bits): 2048
              Certificate Start Date: Fri Apr 30 14:14:46 2010
         Certificate Expiration Date: Sat Apr 30 14:14:46 2011
              Public Key Certificate: -----BEGIN CERTIFICATE-----

MIIDfTCCAmWgAwIBAwIBADANBgkqhkiG9w0BAQsFADBgMRQwEgYDVQQDEwtsYWIu

YWJjLmNvbTELMAkGA1UEBhMCVVMxCTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYD

VQQKEwAxCTAHBgNVBAsTADEPMA0GCSqGSIb3DQEJARYAMB4XDTEwMDQzMDE4MTQ0

BgNVHQ8BAf8EBAMCAQYwHQYDVR0OBBYEFCVG7dYGe51akE14ecaCdL+LOAxUMA0G

CSqGSIb3DQEBCwUAA4IBAQBJlE51pkDY3ZpsSrQeMOoWLteIR+1H0wKZOM1Bhy6Q

+gsE3XEtnN07AE4npjIT0eVP0nI9QIJAbP0uPKaCGAVBSBMoM2mOwbfswI7aJoEh

+XuEoNr0GOz+mltnfhgvl1fT6Ms+xzd3LGZYQTworus2
                                        -----END CERTIFICATE-----
        Country Name (2 letter code): US
  State or Province Name (full name): California
           Locality Name (e.g. city): Sunnyvale
    Organization Name (e.g. company): example
    Organization Unit (e.g. section): IT
        Email Address (Contact Name):  web@example.com
                            Protocol: SSL
                    Hashing Function: SHA256
```

security certificate show-user-installed

Display user installed certificates

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about the user installed digital certificates. Some details are displayed only when you use the command with the -instance parameter. In systems upgraded to ONTAP 9.4 or later, existing ONTAP generated certificates will also be shown as part of this command.

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.

[-vserver <Vserver Name>] - Name of Vserver

Selects the Vserver whose digital certificates you want to display.

[-common-name <FQDN or Custom Common Name>] - FQDN or Custom Common Name

Selects the certificates that match this parameter value.

[-serial <text>] - Serial Number of Certificate

Selects the certificates that match this parameter value.

[-ca <text>] - Certificate Authority

Selects the certificates that match this parameter value.

[-type <type of certificate>] - Type of Certificate

Selects the certificates that match this parameter value.

[-subtype <kmip-cert>] - (DEPRECATED)-Certificate Subtype

This parameter has been deprecated in ONTAP 9.6 and may be removed in a future release of ONTAP.

```
Selects the certificate subtype that matches the specified value. The
valid values are as follows:
```

- *kmip-cert* - this is a Key Management Interoperability Protocol (KMIP) certificate

[-cert-name <text>] - Unique Certificate Name

This specifies the system's internal identifier for the certificate. It is unique within a Vserver.

[-size <size of requested certificate in bits>] - Size of Requested Certificate in Bits

Selects the certificates that match this parameter value.

[-start <Date>] - Certificate Start Date

Selects the certificates that match this parameter value.

[-expiration <Date>] - Certificate Expiration Date

Selects the certificates that match this parameter value.

[-public-cert <certificate>] - Public Key Certificate

Selects the certificates that match this parameter value.

[-country <text>] - Country Name

Selects the certificates that match this parameter value.

[-state <text>] - State or Province Name

Selects the certificates that match this parameter value.

[-locality <text>] - Locality Name

Selects the certificates that match this parameter value.

[-organization <text>] - Organization Name

Selects the certificates that match this parameter value.

[-unit <text>] - Organization Unit

Selects the certificates that match this parameter value.

[-email-addr <mail address>] - Contact Administrator's Email Address

Selects the certificates that match this parameter value.

[-protocol <protocol>] - Protocol

Selects the certificates that match this parameter value.

[-hash-function <hashing function>] - Hashing Function

Selects the certificates that match this parameter value.

[-self-signed {true|false}] - Self-Signed Certificate

Selects the certificates that match this parameter value.

[-is-root {true|false}] - Is Root CA Certificate?

Selects the certificates that match this parameter value.

[-authority-key-identifier <text>] - Authority Key Identifier

Selects the certificates that match this parameter value.

[-subject-key-identifier <text>] - Subject Key Identifier

Selects the certificates that match this parameter value.

[-rfc822-name <mail address>,…] - Email Address SAN

Selects the certificates that match this parameter value.

[-uri <text>,…] - URI SAN

Selects the certificates that match this parameter value.

[-dns-name <text>,…] - DNS Name SAN

Selects the certificates that match this parameter value.

[-ipaddr <IP Address>,…] - IP Address SAN

Selects the certificates that match this parameter value.

Examples

The examples below display information about user installed digital certificates.

```
cluster1::> security certificate show-user-installed

Vserver    Serial Number   Certificate Name                          Type
---------- --------------- ------------------------------------------
---------
vs0        4F4E4D7B         www.example.com         server
    Certificate Authority:  www.example.com
          Expiration Date: Thu Feb 28 16:08:28 2013
```

```
cluster1::> security certificate show-user-installed -instance
                             Vserver: vs0
                    Certificate Name:  www.example.com
          FQDN or Custom Common Name:  www.example.com
        Serial Number of Certificate: 4F4E4D7B
               Certificate Authority:  www.example.com
                 Type of Certificate: server
 Size of Requested Certificate(bits): 2048
              Certificate Start Date: Fri Apr 30 14:14:46 2010
         Certificate Expiration Date: Sat Apr 30 14:14:46 2011
              Public Key Certificate: -----BEGIN CERTIFICATE-----

MIIDfTCCAmWgAwIBAwIBADANBgkqhkiG9w0BAQsFADBgMRQwEgYDVQQDEwtsYWIu

YWJjLmNvbTELMAkGA1UEBhMCVVMxCTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYD

VQQKEwAxCTAHBgNVBAsTADEPMA0GCSqGSIb3DQEJARYAMB4XDTEwMDQzMDE4MTQ0

BgNVHQ8BAf8EBAMCAQYwHQYDVR0OBBYEFCVG7dYGe51akE14ecaCdL+LOAxUMA0G

CSqGSIb3DQEBCwUAA4IBAQBJlE51pkDY3ZpsSrQeMOoWLteIR+1H0wKZOM1Bhy6Q

+gsE3XEtnN07AE4npjIT0eVP0nI9QIJAbP0uPKaCGAVBSBMoM2mOwbfswI7aJoEh

+XuEoNr0GOz+mltnfhgvl1fT6Ms+xzd3LGZYQTworus2
                                        -----END CERTIFICATE-----
        Country Name (2 letter code): US
  State or Province Name (full name): California
           Locality Name (e.g. city): Sunnyvale
    Organization Name (e.g. company): example
    Organization Unit (e.g. section): IT
        Email Address (Contact Name):  web@example.com
                            Protocol: SSL
                    Hashing Function: SHA256
```

security certificate show

Display Installed Digital Certificates

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about the installed digital certificates. Some details are displayed only when you use the command with the -instance parameter.

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.

[-vserver <Vserver Name>] - Name of Vserver

Selects the Vserver whose digital certificates you want to display.

[-common-name <FQDN or Custom Common Name>] - FQDN or Custom Common Name

Selects the certificates that match this parameter value.

[-serial <text>] - Serial Number of Certificate

Selects the certificates that match this parameter value.

[-ca <text>] - Certificate Authority

Selects the certificates that match this parameter value.

[-type <type of certificate>] - Type of Certificate

Selects the certificates that match this parameter value.

[-subtype <kmip-cert>] - (DEPRECATED)-Certificate Subtype

This parameter has been deprecated in ONTAP 9.6 and may be removed in a future release of ONTAP.

```
Selects the certificate subtype that matches the specified value. The
valid values are as follows:
```

- *kmip-cert* - this is a Key Management Interoperability Protocol (KMIP) certificate

[-cert-name <text>] - Unique Certificate Name

This specifies the system's internal identifier for the certificate. It is unique within a Vserver.

[-size <size of requested certificate in bits>] - Size of Requested Certificate in Bits

Selects the certificates that match this parameter value.

[-start <Date>] - Certificate Start Date

Selects the certificates that match this parameter value.

[-expiration <Date>] - Certificate Expiration Date

Selects the certificates that match this parameter value.

[-public-cert <certificate>] - Public Key Certificate

Selects the certificates that match this parameter value.

[-country <text>] - Country Name

Selects the certificates that match this parameter value.

[-state <text>] - State or Province Name

Selects the certificates that match this parameter value.

[-locality <text>] - Locality Name

Selects the certificates that match this parameter value.

[-organization <text>] - Organization Name

Selects the certificates that match this parameter value.

[-unit <text>] - Organization Unit

Selects the certificates that match this parameter value.

[-email-addr <mail address>] - Contact Administrator's Email Address

Selects the certificates that match this parameter value.

[-protocol <protocol>] - Protocol

Selects the certificates that match this parameter value.

[-hash-function <hashing function>] - Hashing Function

Selects the certificates that match this parameter value.

[-self-signed {true|false}] - Self-Signed Certificate

Selects the certificates that match this parameter value.

[-is-root {true|false}] - Is Root CA Certificate?

Selects the certificates that match this parameter value.

[-authority-key-identifier <text>] - Authority Key Identifier

Selects the certificates that match this parameter value.

[-subject-key-identifier <text>] - Subject Key Identifier

Selects the certificates that match this parameter value.

[-rfc822-name <mail address>,…] - Email Address SAN

Selects the certificates that match this parameter value.

[-uri <text>,…] - URI SAN

Selects the certificates that match this parameter value.

[-dns-name <text>,…] - DNS Name SAN

Selects the certificates that match this parameter value.

[-ipaddr <IP Address>,…] - IP Address SAN

Selects the certificates that match this parameter value.

Examples

The examples below display information about digital certificates.

```
cluster1::> security certificate show

Vserver    Serial Number   Certificate Name                          Type
---------- --------------- ------------------------------------------
---------
vs0        4F4E4D7B         www.example.com
server
    Certificate Authority:  www.example.com
          Expiration Date: Thu Feb 28 16:08:28 2013
```

```
cluster1::> security certificate show -instance
                             Vserver: vs0
                    Certificate Name:  www.example.com
          FQDN or Custom Common Name:  www.example.com
        Serial Number of Certificate: 4F4E4D7B
               Certificate Authority:  www.example.com
                 Type of Certificate: server
 Size of Requested Certificate(bits): 2048
              Certificate Start Date: Fri Apr 30 14:14:46 2010
         Certificate Expiration Date: Sat Apr 30 14:14:46 2011
              Public Key Certificate: -----BEGIN CERTIFICATE-----

MIIDfTCCAmWgAwIBAwIBADANBgkqhkiG9w0BAQsFADBgMRQwEgYDVQQDEwtsYWIu

YWJjLmNvbTELMAkGA1UEBhMCVVMxCTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYD

VQQKEwAxCTAHBgNVBAsTADEPMA0GCSqGSIb3DQEJARYAMB4XDTEwMDQzMDE4MTQ0

BgNVHQ8BAf8EBAMCAQYwHQYDVR0OBBYEFCVG7dYGe51akE14ecaCdL+LOAxUMA0G

CSqGSIb3DQEBCwUAA4IBAQBJlE51pkDY3ZpsSrQeMOoWLteIR+1H0wKZOM1Bhy6Q

+gsE3XEtnN07AE4npjIT0eVP0nI9QIJAbP0uPKaCGAVBSBMoM2mOwbfswI7aJoEh

+XuEoNr0GOz+mltnfhgvl1fT6Ms+xzd3LGZYQTworus2
                                        -----END CERTIFICATE-----
        Country Name (2 letter code): US
  State or Province Name (full name): California
           Locality Name (e.g. city): Sunnyvale
    Organization Name (e.g. company): example
    Organization Unit (e.g. section): IT
        Email Address (Contact Name):  web@example.com
                            Protocol: SSL
                    Hashing Function: SHA256
```

security certificate sign

Sign a Digital Certificate using Self-Signed Root CA

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command signs a digital certificate signing request and generates a certificate using a Self-Signed Root CA certificate in either PEM or PKCS12 format. You can use the security certificate generate-csr command to generate a digital certificate signing request.

Parameters

`-vserver <Vserver Name>` - Name of Vserver

This specifies the name of the Vserver on which the signed certificate will exist.

`-ca <text>` - Certificate Authority to Sign

This specifies the name of the Certificate Authority that will sign the certificate.

`-ca-serial <text>` - Serial Number of CA Certificate

This specifies the serial number of the Certificate Authority that will sign the certificate.

`[-expire-days <integer>]` - Number of Days until Expiration

This specifies the number of days until the signed certificate expires. The default value is 365 days. Possible values are between *1* and *3652* .

`[-format <certificate format>]` - Certificate Format

This specifies the format of signed certificate. The default value is PEM. Possible values include *PEM* and *PKCS12* .

`[-destination {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}]` - Where to Send File

This specifies the destination to upload the signed certificate. This option can only be used when the format is PKCS12.

`[-hash-function <hashing function>]` - Hashing Function

This specifies the cryptographic hashing function for the self-signed certificate. The default value is SHA256. Possible values include *SHA224* , *SHA256* , *SHA384* , and *SHA512* .

Examples

This example signs a digital certificate for a Vserver named vs0 using a Certificate Authority certificate that has a ca of *www.ca.com* and a ca-serial of 4F4EB629 in PEM format using the SHA256 hashing function.

```
cluster1::> security certificate sign -vserver vs0 -ca  www.ca.com -ca
-serial 4F4EB629 -expire-days 36 -format PEM -hash-function SHA256

Enter certificate signing request (CSR): Press <Enter> when done
-----BEGIN CERTIFICATE REQUEST-----
MIIBGjCBxQIBADBgMRQwEgYDVQQDEwtleGFtcGxlLmNvbTELMAkGA1UEBhMCVVMx
CTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYDVQQKEwAxCTAHBgNVBAsTADEPMA0G
CSqGSIb3DQEJARYAMFwwDQYJKoZIhvcNAQEBBQADSwAwSAJBAPXFanNoJApT1nzS
xOcxixqImRRGZCR7tVmTYyqPSuTvfhVtwDJbmXuj6U3a1woUsb13wfEvQnHVFNci
2ninsJ8CAwEAAaAAMA0GCSqGSIb3DQEBCwUAA0EA6EagLfso5+4g+ejiRKKTUPQO
UqOUEoKuvxhOvPC2w7b//fNSFsFHvXloqEOhYECn/NX9h8mbphCoM5YZ4OfnKw==
-----END CERTIFICATE REQUEST-----
Signed Certificate: :
-----BEGIN CERTIFICATE-----
MIICwDCCAaigAwIBAgIET1oskDANBgkqhkiG9w0BAQsFADBdMREwDwYDVQQDEwh2
czAuY2VydDELMAkGA1UEBhMCVVMxCTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYD
VQQKEwAxCTAHBgNVBAsTADEPMA0GCSqGSIb3DQEJARYAMB4XDTEyMDMwOTE2MTUx
M1oXDTEyMDQxNDE2MTUxM1owYDEUMBIGA1UEAxMLZXhhbXBsZS5jb20xCzAJBgNV
BAYTAlVTMQkwBwYDVQQIEwAxCTAHBgNVBAcTADEJMAcGA1UEChMAMQkwBwYDVQQL
EwAxDzANBgkqhkiG9w0BCQEWADBcMA0GCSqGSIb3DQEBAQUAA0sAMEgCQQD1xWpz
-----END CERTIFICATE-----
```

This example signs and exports a digital certificate to destination ftp://10.98.1.1//u/sam/sign.pfx for a Vserver named vs0 using a Certificate Authority certificate that expires in 36 days and has a ca value of *www.ca.com* and a ca-serial value of 4F4EB629 in PKCS12 format by the SHA384 hashing function.

```
cluster1::> security certificate sign -vserver vs0 -ca www.ca.com -ca
-serial 4F4EB629
-expire-days 36 -format PKCS12 -destination
ftp://10.98.1.1//u/sam/sign.pfx -hash-function SHA384

Enter certificate signing request (CSR): Press <Enter> when done
-----BEGIN CERTIFICATE REQUEST-----
MIIBGjCBxQIBADBgMRQwEgYDVQQDEwtleGFtcGxlLmNvbTELMAkGA1UEBhMCVVMx
CTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYDVQQKEwAxCTAHBgNVBAsTADEPMA0G
CSqGSIb3DQEJARYAMFwwDQYJKoZIhvcNAQEBBQADSwAwSAJBAPXFanNoJApT1nzS
xOcxixqImRRGZCR7tVmTYyqPSuTvfhVtwDJbmXuj6U3a1woUsb13wfEvQnHVFNci
2ninsJ8CAwEAAaAAMA0GCSqGSIb3DQEBCwUAA0EA6EagLfso5+4g+ejiRKKTUPQO
UqOUEoKuvxhOvPC2w7b//fNSFsFHvXloqEOhYECn/NX9h8mbphCoM5YZ4OfnKw==
-----END CERTIFICATE REQUEST-----
Signed Certificate: :
-----BEGIN CERTIFICATE-----
MIICwDCCAaigAwIBAgIET1ot8jANBgkqhkiG9w0BAQsFADBdMREwDwYDVQQDEwh2
czAuY2VydDELMAkGA1UEBhMCVVMxCTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYD
VQQKEwAxCTAHBgNVBAsTADEPMA0GCSqGSIb3DQEJARYAMB4XDTEyMDMwOTE2MjEw
NloXDTEyMDQxNDE2MjEwNlowYDEUMBIGA1UEAxMLZXhhbXBsZS5jb20xCzAJBgNV
BAYTAlVTMQkwBwYDVQQIEwAxCTAHBgNVBAcTADEJMAcGA1UEChMAMQkwBwYDVQQL
EwAxDzANBgkqhkiG9w0BCQEWADBcMA0GCSqGSIb3DQEBAQUAA0sAMEgCQQD1xWpz
oarXHSyDzv3T5QIxBGRJ0ACtgdjJuqtuAdmnKvKfLS1o4C90
-----END CERTIFICATE-----
Enter private key: Press <Enter> when done
-----BEGIN RSA PRIVATE KEY-----
MIIBOwIBAAJBAPXFanNoJApT1nzSxOcxixqImRRGZCR7tVmTYyqPSuTvfhVtwDJb
mXuj6U3a1woUsb13wfEvQnHVFNci2ninsJ8CAwEAAQJAWt2AO+bW3FKezEuIrQlu
KoMyRYK455wtMk8BrOyJfhYsB20B28eifjJvRWdTOBEav99M7cEzgPv+p5kaZTTM
gQIhAPsp+j1hrUXSRj979LIJJY0sNez397i7ViFXWQScx/ehAiEA+oDbOooWlVvu
xj4aitxVBu6ByVckYU8LbsfeRNsZwD8CIQCbZ1/ENvmlJ/P7N9Exj2NCtEYxd0Q5
cwBZ5NfZeMBpwQIhAPk0KWQSLadGfsKO077itF+h9FGFNHbtuNTrVq4vPW3nAiAA
peMBQgEv28y2r8D4dkYzxcXmjzJluUSZSZ9c/wS6fA==
-----END RSA PRIVATE KEY-----
Enter a password for pkcs12 file:
Enter it again:

Enter User for Destination URI: sam
Enter Password:
```

Related Links

- security certificate generate-csr

security certificate ca-issued revoke

Revoke a Digital Certificate

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command revokes a digital certificate signed by a Self-Signed Root CA.

Parameters

`-vserver <Vserver Name>` - Name of Vserver

This specifies the name of the Vserver on which the certificate is stored.

`-serial <text>` - Serial Number of Certificate

This specifies the serial number of the certificate.

`-ca <text>` - Certificate Authority

This specifies the name of the Certificate Authority whose certificate will be revoked.

`-ca-serial <text>` - Serial Number of CA Certificate

This specifies the serial number of Certificate Authority.

`[-common-name <FQDN or Custom Common Name>]` - FQDN or Custom Common Name

This specifies a fully qualified domain name (FQDN) or custom common name or the name of a person. This field is optional if ca-serial is specified.

Examples

This example revokes a signed digital certificate for a Vserver named vs0 with serial as 4F5A2DF2 for a Certificate Authority certificate that has a ca of *`www.ca.com`* and a ca-serial of 4F4EB629.

```
cluster1::> security certificate ca-issued revoke -vserver vs0 -serial
4F5A2DF2 -ca  www.ca.com  -ca-serial 4F4EB629
```

security certificate ca-issued show

Display CA-Issued Digital Certificates

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the following information about the digital certificates issued by the self-signed root-ca:

- Vserver
- Serial number of certificate
- FQDN or custom common name or the name of a person

- Serial number of CA certificate
- Status (`active, revoked`)
- Certificate Authority
- Expiration date
- Revocation date

To display more details, run the command with the `-instance` parameter. This will add the following information:

- Country name
- State or province name
- Locality name
- Organization name
- Organization unit
- Contact administrator's email address

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.

[-vserver <Vserver Name>] - Name of Vserver

Selects the certificates that match this parameter value.

[-serial <text>] - Serial Number of Certificate

Selects the certificates that match this parameter value.

[-ca <text>] - Certificate Authority

Selects the certificates that match this parameter value.

[-ca-serial <text>] - Serial Number of CA Certificate

Selects the certificates that match this parameter value.

[-common-name <FQDN or Custom Common Name>] - FQDN or Custom Common Name

Selects the certificates that match this parameter value.

[-status <status of certificate>] - Status of Certificate

Selects the certificates that match this parameter value. Possible values include active and revoked.

[-expiration <Date>] - Certificate Expiration Date

Selects the certificates that match this parameter value.

[-revocation <Date>] - Certificate Revocation Date

Selects the certificates that match this parameter value.

[-country <text>] - Country Name (2 letter code)

Selects the certificates that match this parameter value.

[-state <text>] - State or Province Name (full name)

Selects the certificates that match this parameter value.

[-locality <text>] - Locality Name (e.g. city)

Selects the certificates that match this parameter value.

[-organization <text>] - Organization Name (e.g. company)

Selects the certificates that match this parameter value.

[-unit <text>] - Organization Unit (e.g. section)

Selects the certificates that match this parameter value.

[-email-addr <mail address>] - Email Address (Contact Name)

Selects the certificates that match this parameter value.

Examples

The examples below display information about CA issued digital certificates.

```
cluster1::> security certificate ca-issued show
Serial Number of
Vserver    Serial Number   Common Name                 CA's Certificate
Status
---------- --------------- --------------------------- ----------------
-------
vs0        4F5A2C90         example.com                  4F4EB629
active
    Certificate Authority: vs0.cert
          Expiration Date: Sat Apr 14 16:15:13 2012
          Revocation Date: -

vs0        4F5A2DF2         example.com                  4F4EB629
revoked
    Certificate Authority: vs0.cert
          Expiration Date: Sat Apr 14 16:21:06 2012
          Revocation Date: Fri Mar 09 17:08:30 2012

2 entries were displayed.
```

```
cluster1::> security certificate ca-issued show -instance
Vserver: vs0
         Serial Number of Certificate: 4F5A2C90
                Certificate Authority: vs0.cert
      Serial Number of CA Certificate: 4F4EB629
           FQDN or Custom Common Name:  example.com
                Status of Certificate: active
          Certificate Expiration Date: Sat Apr 14 16:15:13 2012
          Certificate Revocation Date: -
         Country Name (2 letter code): US
   State or Province Name (full name): California
            Locality Name (e.g. city): Sunnyvale
     Organization Name (e.g. company): example
     Organization Unit (e.g. section): IT
         Email Address (Contact Name):  web@example.com
```

security certificate config modify

Modify the certificate management configurations

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command modifies the certificate management configuration information for the cluster.

Parameters

`[-min-security-strength <bits of security strength>]` - Minimum Security Strength

Use this parameter to modify the allowed minimum security strength for certificates. The security bits mapping to RSA and ECDSA key length are as follows:

```
            Security Bits   Asymmetric Key Length   Elliptic Curve Key
Length
            112             2048                    224
            128             3072                    256
            192             4096                    384
```

```
FIPS supported values are restricted to 112 and 128.
```

+
NOTE: This does not affect root CA certificates.

+

[-expiration-warn-threshold <integer>] - Minimum Days to EMS for Expiring Certificates

Use this parameter to modify the number of days prior to certificate expiration the system sends a warning EMS event.

Examples

The following example modifies the minimum security strength allowed for certificates.

```
cluster-1::> security certificate config modify -min-security-strength 192
```

security certificate config show

Displays the certificate management configurations

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays the certificate management configuration information for the cluster.

"min-security-strength" - If you specify this parameter, the command displays the minimum allowed security strength for certificates.

"expiration-warn-threshold" - If you specify this parameter, the command displays the minimum number of days before expiration date configured for event management system (EMS) notification of expiring certificates.

Examples

The following example lists minimum security strength certificate management configuration.

```
cluster-1::> security certificate config show -fields min-security-
strength

Minimum Security Strength
 -------------------------

112
```

security certificate truststore check

Initiate a TLS connection and identify the root CA certificate

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command allows the user to check if the node can use the installed set of CA certificates to establish a secure connection with the specified server. If the connection attempt fails, the system reports which expected certificates are missing. If the attempt succeeds, the system displays details of the certificates used.

Parameters

`-vserver <Vserver Name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver that needs the connectivity check.

`-server <Hostname and Port>` - Server Name (privilege: advanced)

Use this parameter to specify the server to establish a connection with and look up the required CA certificate.

Examples

The following example demonstrates a missing CA certificate:

```
cluster1::*> security certificate truststore check -vserver cluster1
-server example.com:443

Error: command failed: Missing certificate with subject name: "CN =
ExampleRoot, C = US"
```

The following example demonstrates the required certificate being present:

```
cluster1::*> security certificate truststore check -server example.com:443

CA certificate with cert-name "ExampleRoot" is already installed in the
truststore. Use "security certificate show -cert-name ExampleRoot" to see
the details of the CA certificate.
```

security certificate truststore clear

Clear the default root certificates from truststore

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security certificate truststore clear` command deletes the pre-installed certificates of the type 'server-ca'. If you delete these certificates, some of the applications performing SSL communication can fail.

Examples

The following example removes the default certificate bundle:

```
cluster1::> security certificate truststore clear
```

security certificate truststore load

Load the default root certificates to truststore

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security certificate truststore load` command installs default root certificates in the certificate table of type 'server-ca'. These are the certificates required to validate any incoming server certificate during the SSL handshake process. Note: This command only installs PEM formatted certificates.

Parameters

{ `[-path <text>]` - File to load PEM certificates from (privilege: advanced)

This specifies the path to the PEM formatted certificate bundle. This optional parameter can have an empty value (the default).

| `[-uri <text>]` - URL to download PEM certificates from (privilege: advanced) }

This specifies the URL from which to download the PEM formatted certificate bundle.

`[-ontap-version <ontap_version>]` - Certificates from specific ONTAP version (privilege: advanced)

This specifies the ONTAP version in which the certificates were introduced. Only those certificates will be loaded. This optional parameter can have an empty value (the default) which indicates that no filtering on version is done.

Examples

The following example installs the default certificate bundle:

```
cluster1::> security certificate truststore load
```

security config commands

security config modify

Modify Security Configuration Options

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security config modify` command modifies the existing cluster-wide security configuration. If you enable FIPS-compliant mode, the cluster will automatically select only compliant TLS protocols (currently TLSv1.3 and TLSv1.2). Non-compliant protocols are not enabled when FIPS-compliant mode is disabled. Use the `-supported-protocols` parameter to include or exclude TLS protocols independently from the FIPS mode. All protocols at or above the lowest version specified will be enabled, even those not explicitly specified. By default, FIPS mode is disabled, and ONTAP supports the TLSv1.3 and TLSv1.2 protocols. For backward compatibility, ONTAP supports adding SSLv3 and TLSv1 to the supported-protocols list when FIPS mode is

disabled. Use the `-supported-cipher-suites` parameter to control which TLS cipher suites are permitted by the system. By default the supported-cipher-suites setting is *TLS_RSA_WITH_AES_128_CCM, TLS_RSA_WITH_AES_128_CCM_8, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CCM, TLS_RSA_WITH_AES_256_CCM_8, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_ARIA_128_GCM_SHA256, TLS_RSA_WITH_ARIA_256_GCM_SHA384, TLS_RSA_WITH_CAMELLIA_128_CBC_SHA, TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256, TLS_RSA_WITH_CAMELLIA_256_CBC_SHA, TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_ARIA_128_GCM_SHA256, TLS_DHE_DSS_WITH_ARIA_256_GCM_SHA384, TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA, TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA256, TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA, TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA256, TLS_DHE_PSK_WITH_AES_128_CBC_SHA, TLS_DHE_PSK_WITH_AES_128_CBC_SHA256, TLS_DHE_PSK_WITH_AES_128_CCM, TLS_PSK_DHE_WITH_AES_128_CCM_8, TLS_DHE_PSK_WITH_AES_128_GCM_SHA256, TLS_DHE_PSK_WITH_AES_256_CBC_SHA, TLS_DHE_PSK_WITH_AES_256_CBC_SHA384, TLS_DHE_PSK_WITH_AES_256_CCM, TLS_PSK_DHE_WITH_AES_256_CCM_8, TLS_DHE_PSK_WITH_AES_256_GCM_SHA384, TLS_DHE_PSK_WITH_ARIA_128_GCM_SHA256, TLS_DHE_PSK_WITH_ARIA_256_GCM_SHA384, TLS_DHE_PSK_WITH_CAMELLIA_128_CBC_SHA256, TLS_DHE_PSK_WITH_CAMELLIA_256_CBC_SHA384, TLS_DHE_PSK_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_128_CCM, TLS_DHE_RSA_WITH_AES_128_CCM_8, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CCM, TLS_DHE_RSA_WITH_AES_256_CCM_8, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_ARIA_128_GCM_SHA256, TLS_DHE_RSA_WITH_ARIA_256_GCM_SHA384, TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA, TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256, TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA, TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDHE_RSA_WITH_ARIA_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_ARIA_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CCM, TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CCM, TLS_ECDHE_ECDSA_WITH_AES_256_CCM_8, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_ARIA_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_ARIA_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_CAMELLIA_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_CAMELLIA_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA, TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA256, TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA, TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA384, TLS_ECDHE_PSK_WITH_CAMELLIA_128_CBC_SHA256, TLS_ECDHE_PSK_WITH_CAMELLIA_256_CBC_SHA384, TLS_ECDHE_PSK_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,*

TLS_ECDHE_RSA_WITH_CAMELLIA_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_PSK_WITH_AES_128_CBC_SHA, TLS_PSK_WITH_AES_128_CBC_SHA256, TLS_PSK_WITH_AES_128_CCM, TLS_PSK_WITH_AES_128_CCM_8, TLS_PSK_WITH_AES_128_GCM_SHA256, TLS_PSK_WITH_AES_256_CBC_SHA, TLS_PSK_WITH_AES_256_CBC_SHA384, TLS_PSK_WITH_AES_256_CCM, TLS_PSK_WITH_AES_256_CCM_8, TLS_PSK_WITH_AES_256_GCM_SHA384, TLS_PSK_WITH_ARIA_128_GCM_SHA256, TLS_PSK_WITH_ARIA_256_GCM_SHA384, TLS_PSK_WITH_CAMELLIA_128_CBC_SHA256, TLS_PSK_WITH_CAMELLIA_256_CBC_SHA384, TLS_PSK_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_PSK_WITH_AES_128_CBC_SHA, TLS_RSA_PSK_WITH_AES_128_CBC_SHA256, TLS_RSA_PSK_WITH_AES_128_GCM_SHA256, TLS_RSA_PSK_WITH_AES_256_CBC_SHA, TLS_RSA_PSK_WITH_AES_256_CBC_SHA384, TLS_RSA_PSK_WITH_AES_256_GCM_SHA384, TLS_RSA_PSK_WITH_ARIA_128_GCM_SHA256, TLS_RSA_PSK_WITH_ARIA_256_GCM_SHA384, TLS_RSA_PSK_WITH_CAMELLIA_128_CBC_SHA256, TLS_RSA_PSK_WITH_CAMELLIA_256_CBC_SHA384, TLS_RSA_PSK_WITH_CHACHA20_POLY1305_SHA256, TLS_SRP_SHA_WITH_AES_128_CBC_SHA, TLS_SRP_SHA_WITH_AES_256_CBC_SHA, TLS_SRP_SHA_DSS_WITH_AES_128_CBC_SHA, TLS_SRP_SHA_DSS_WITH_AES_256_CBC_SHA, TLS_SRP_SHA_RSA_WITH_AES_128_CBC_SHA, TLS_SRP_SHA_RSA_WITH_AES_256_CBC_SHA, TLS_AES_128_GCM_SHA256, TLS_AES_256_GCM_SHA384, TLS_CHACHA20_POLY1305_SHA256. Select a cipher suite which is available with the corresponding selected protocol. An invalid configuration may cause some functionality to fail to operate properly. Valid values for supported-cipher-suites are listed at "https://www.iana.org/assignments/tls-parameters/tls-parameters.xhtml" published by IANA.

Parameters

`-interface <SSL>` - (DEPRECATED)-FIPS-Compliant Interface (privilege: advanced)

This parameter has been deprecated in ONTAP 9.8 and may be removed in a future release of ONTAP.

```
Selects the FIPS-compliant interface. The only valid value is ``_SSL_``
.
```

`[-is-fips-enabled {true|false}]` - FIPS Mode (privilege: advanced)

Enables or disables FIPS-compliant mode for the entire cluster. Default is *`false`*.

`[-supported-protocols {TLSv1.3|TLSv1.2|TLSv1.1|TLSv1|SSLv3}]` - Supported Protocols (privilege: advanced)

Selects the supported protocols for the selected interface. Default is *`TLSv1.3,TLSv1.2`*

`[-supported-ciphers <Cipher String>]` - (DEPRECATED)-Supported Ciphers (privilege: advanced)

This parameter has been deprecated in ONTAP 9.8 and may be removed in a future release of ONTAP. Use the supported-ciphers-suites parameter instead.

```
Selects the supported cipher suites for the selected interface. Default
is ``_ALL:!LOW:!aNULL:!EXP:!eNULL_`` .
```

`[-supported-cipher-suites <Cipher String>,…]` - Supported Cipher Suites (privilege: advanced)

Selects the supported cipher suites for the selected interface.

Examples

The following command enables FIPS mode in the cluster. (Default setting for FIPS mode is *`false`*)

```
cluster1::> security config modify * -is-fips-enabled true
```

The following command limits the supported protocols to just TLSv1.3 in the cluster. (Default setting for supported protocols is *`TLSv1.3,TLSv1.2`*)

```
cluster1::*> security config modify * -supported-protocols TLSv1.3
```

The following command limits the supported cipher suites in the cluster to the listed ciphers.

```
cluster1::*> security config modify * -supported-cipher-suites
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,TLS_AES_256_GCM_SHA384
```

security config show

Display Security Configuration Options

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security config show` command displays the security configurations of the cluster in advanced privilege mode.

Default values are as follows:

- SSL FIPS mode: disabled
- Supported protocols: TLSv1.3,TLSv1.2
- Supported cipher suites: All suites for the listed protocols except those that have no authentication, low encryption strength (less than 56 bits), or use 3DES or static DH key exchange.

Enabling FIPS mode will cause the entire cluster to use FIPS-compliant crypto operations only.

Use the security config modify command to change the protocols and cipher suites that the cluster will support.

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.

[-interface <SSL>] - (DEPRECATED)-FIPS-Compliant Interface (privilege: advanced)

> This parameter has been deprecated in ONTAP 9.8 and may be removed in a future release of ONTAP. As there only ever existed one valid value for this parameter, filtering on it has never altered the results.

```
Displays configurations that match the specified value for the
interface.
```

[-is-fips-enabled {true|false}] - FIPS Mode (privilege: advanced)

Display configurations that match the specified value for FIPS mode.

[-supported-protocols {TLSv1.3|TLSv1.2|TLSv1.1|TLSv1|SSLv3}] - Supported Protocols (privilege: advanced)

Displays configurations that match the specified protocols.

[-supported-ciphers <Cipher String>] - (DEPRECATED)-Supported Ciphers (privilege: advanced)

> This parameter has been deprecated in ONTAP 9.8 and may be removed in a future release of ONTAP. Use the supported-cipher-suites parameter instead.

```
Displays the configurations that match the specified supported ciphers.
```

[-supported-cipher-suites <Cipher String>,…] - Supported Cipher Suites (privilege: advanced)

Displays the configurations that match the specified supported cipher suites.

Examples

The following example shows the default security configurations for a cluster.

```
cluster1::> security config show
Cluster    Supported
FIPS Mode  Protocols Supported Cipher Suites
---------- ---------
------------------------------------------------------------
false      TLSv1.3,  TLS_RSA_WITH_AES_128_CCM, TLS_RSA_WITH_AES_128_CCM_8,
```

```
            TLSv1.2,  TLS_RSA_WITH_AES_128_GCM_SHA256,
                      TLS_RSA_WITH_AES_128_CBC_SHA,
                      TLS_RSA_WITH_AES_128_CBC_SHA256,
TLS_RSA_WITH_AES_256_CCM,
                      TLS_RSA_WITH_AES_256_CCM_8,
                      TLS_RSA_WITH_AES_256_GCM_SHA384,
                      TLS_RSA_WITH_AES_256_CBC_SHA,
                      TLS_RSA_WITH_AES_256_CBC_SHA256,
                      TLS_RSA_WITH_ARIA_128_GCM_SHA256,
                      TLS_RSA_WITH_ARIA_256_GCM_SHA384,
                      TLS_RSA_WITH_CAMELLIA_128_CBC_SHA,
                      TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256,
                      TLS_RSA_WITH_CAMELLIA_256_CBC_SHA,
                      TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256,
                      TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,
                      TLS_DHE_DSS_WITH_AES_128_CBC_SHA,
                      TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,
                      TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,
                      TLS_DHE_DSS_WITH_AES_256_CBC_SHA,
                      TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,
                      TLS_DHE_DSS_WITH_ARIA_128_GCM_SHA256,
                      TLS_DHE_DSS_WITH_ARIA_256_GCM_SHA384,
                      TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA,
                      TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA256,
                      TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA,
                      TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA256,
                      TLS_DHE_PSK_WITH_AES_128_CBC_SHA,
                      TLS_DHE_PSK_WITH_AES_128_CBC_SHA256,
                      TLS_DHE_PSK_WITH_AES_128_CCM,
                      TLS_PSK_DHE_WITH_AES_128_CCM_8,
                      TLS_DHE_PSK_WITH_AES_128_GCM_SHA256,
                      TLS_DHE_PSK_WITH_AES_256_CBC_SHA,
                      TLS_DHE_PSK_WITH_AES_256_CBC_SHA384,
                      TLS_DHE_PSK_WITH_AES_256_CCM,
                      TLS_PSK_DHE_WITH_AES_256_CCM_8,
                      TLS_DHE_PSK_WITH_AES_256_GCM_SHA384,
                      TLS_DHE_PSK_WITH_ARIA_128_GCM_SHA256,
                      TLS_DHE_PSK_WITH_ARIA_256_GCM_SHA384,
                      TLS_DHE_PSK_WITH_CAMELLIA_128_CBC_SHA256,
                      TLS_DHE_PSK_WITH_CAMELLIA_256_CBC_SHA384,
                      TLS_DHE_PSK_WITH_CHACHA20_POLY1305_SHA256,
                      TLS_DHE_RSA_WITH_AES_128_CCM,
                      TLS_DHE_RSA_WITH_AES_128_CCM_8,
                      TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,
                      TLS_DHE_RSA_WITH_AES_128_CBC_SHA,
                      TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,
```

```
                    TLS_DHE_RSA_WITH_AES_256_CCM,
                    TLS_DHE_RSA_WITH_AES_256_CCM_8,
                    TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,
                    TLS_DHE_RSA_WITH_AES_256_CBC_SHA,
                    TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,
                    TLS_DHE_RSA_WITH_ARIA_128_GCM_SHA256,
                    TLS_DHE_RSA_WITH_ARIA_256_GCM_SHA384,
                    TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA,
                    TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256,
                    TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA,
                    TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256,
                    TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256,
                    TLS_ECDHE_RSA_WITH_ARIA_128_GCM_SHA256,
                    TLS_ECDHE_RSA_WITH_ARIA_256_GCM_SHA384,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CCM,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8,
                    TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CCM,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CCM_8,
                    TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,
                    TLS_ECDHE_ECDSA_WITH_ARIA_128_GCM_SHA256,
                    TLS_ECDHE_ECDSA_WITH_ARIA_256_GCM_SHA384,
                    TLS_ECDHE_ECDSA_WITH_CAMELLIA_128_CBC_SHA256,
                    TLS_ECDHE_ECDSA_WITH_CAMELLIA_256_CBC_SHA384,
                    TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256,
                    TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA,
                    TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA256,
                    TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA,
                    TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA384,
                    TLS_ECDHE_PSK_WITH_CAMELLIA_128_CBC_SHA256,
                    TLS_ECDHE_PSK_WITH_CAMELLIA_256_CBC_SHA384,
                    TLS_ECDHE_PSK_WITH_CHACHA20_POLY1305_SHA256,
                    TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
                    TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
                    TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,
                    TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,
                    TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,
                    TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,
                    TLS_ECDHE_RSA_WITH_CAMELLIA_128_CBC_SHA256,
                    TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384,
                    TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256,
                    TLS_PSK_WITH_AES_128_CBC_SHA,
```

```
                      TLS_PSK_WITH_AES_128_CBC_SHA256,
TLS_PSK_WITH_AES_128_CCM,
                      TLS_PSK_WITH_AES_128_CCM_8,
                      TLS_PSK_WITH_AES_128_GCM_SHA256,
                      TLS_PSK_WITH_AES_256_CBC_SHA,
                      TLS_PSK_WITH_AES_256_CBC_SHA384,
TLS_PSK_WITH_AES_256_CCM,
                      TLS_PSK_WITH_AES_256_CCM_8,
                      TLS_PSK_WITH_AES_256_GCM_SHA384,
                      TLS_PSK_WITH_ARIA_128_GCM_SHA256,
                      TLS_PSK_WITH_ARIA_256_GCM_SHA384,
                      TLS_PSK_WITH_CAMELLIA_128_CBC_SHA256,
                      TLS_PSK_WITH_CAMELLIA_256_CBC_SHA384,
                      TLS_PSK_WITH_CHACHA20_POLY1305_SHA256,
                      TLS_RSA_PSK_WITH_AES_128_CBC_SHA,
                      TLS_RSA_PSK_WITH_AES_128_CBC_SHA256,
                      TLS_RSA_PSK_WITH_AES_128_GCM_SHA256,
                      TLS_RSA_PSK_WITH_AES_256_CBC_SHA,
                      TLS_RSA_PSK_WITH_AES_256_CBC_SHA384,
                      TLS_RSA_PSK_WITH_AES_256_GCM_SHA384,
                      TLS_RSA_PSK_WITH_ARIA_128_GCM_SHA256,
                      TLS_RSA_PSK_WITH_ARIA_256_GCM_SHA384,
                      TLS_RSA_PSK_WITH_CAMELLIA_128_CBC_SHA256,
                      TLS_RSA_PSK_WITH_CAMELLIA_256_CBC_SHA384,
                      TLS_RSA_PSK_WITH_CHACHA20_POLY1305_SHA256,
                      TLS_SRP_SHA_WITH_AES_128_CBC_SHA,
                      TLS_SRP_SHA_WITH_AES_256_CBC_SHA,
                      TLS_SRP_SHA_DSS_WITH_AES_128_CBC_SHA,
                      TLS_SRP_SHA_DSS_WITH_AES_256_CBC_SHA,
                      TLS_SRP_SHA_RSA_WITH_AES_128_CBC_SHA,
                      TLS_SRP_SHA_RSA_WITH_AES_256_CBC_SHA
```

The following example shows the security configuration after FIPS mode has been enabled.

```
cluster1::> security config show
Cluster    Supported
FIPS Mode  Protocols Supported Cipher Suites
---------- ---------
------------------------------------------------------------
true       TLSv1.3,  TLS_RSA_WITH_AES_128_CCM, TLS_RSA_WITH_AES_128_CCM_8,
           TLSv1.2,  TLS_RSA_WITH_AES_128_GCM_SHA256,
                     TLS_RSA_WITH_AES_128_CBC_SHA,
                     TLS_RSA_WITH_AES_128_CBC_SHA256,
TLS_RSA_WITH_AES_256_CCM,
                     TLS_RSA_WITH_AES_256_CCM_8,
```

```
                    TLS_RSA_WITH_AES_256_GCM_SHA384,
                    TLS_RSA_WITH_AES_256_CBC_SHA,
                    TLS_RSA_WITH_AES_256_CBC_SHA256,
                    TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,
                    TLS_DHE_DSS_WITH_AES_128_CBC_SHA,
                    TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,
                    TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,
                    TLS_DHE_DSS_WITH_AES_256_CBC_SHA,
                    TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,
                    TLS_DHE_PSK_WITH_AES_128_CBC_SHA,
                    TLS_DHE_PSK_WITH_AES_128_CBC_SHA256,
                    TLS_DHE_PSK_WITH_AES_128_CCM,
                    TLS_PSK_DHE_WITH_AES_128_CCM_8,
                    TLS_DHE_PSK_WITH_AES_128_GCM_SHA256,
                    TLS_DHE_PSK_WITH_AES_256_CBC_SHA,
                    TLS_DHE_PSK_WITH_AES_256_CBC_SHA384,
                    TLS_DHE_PSK_WITH_AES_256_CCM,
                    TLS_PSK_DHE_WITH_AES_256_CCM_8,
                    TLS_DHE_PSK_WITH_AES_256_GCM_SHA384,
                    TLS_DHE_RSA_WITH_AES_128_CCM,
                    TLS_DHE_RSA_WITH_AES_128_CCM_8,
                    TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,
                    TLS_DHE_RSA_WITH_AES_128_CBC_SHA,
                    TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,
                    TLS_DHE_RSA_WITH_AES_256_CCM,
                    TLS_DHE_RSA_WITH_AES_256_CCM_8,
                    TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,
                    TLS_DHE_RSA_WITH_AES_256_CBC_SHA,
                    TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CCM,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8,
                    TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,
                    TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CCM,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CCM_8,
                    TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,
                    TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,
                    TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA,
                    TLS_ECDHE_PSK_WITH_AES_128_CBC_SHA256,
                    TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA,
                    TLS_ECDHE_PSK_WITH_AES_256_CBC_SHA384,
                    TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
                    TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
                    TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,
```

```
                      TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,
                      TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,
                      TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,
                      TLS_PSK_WITH_AES_128_CBC_SHA,
                      TLS_PSK_WITH_AES_128_CBC_SHA256,
TLS_PSK_WITH_AES_128_CCM,
                      TLS_PSK_WITH_AES_128_CCM_8,
                      TLS_PSK_WITH_AES_128_GCM_SHA256,
                      TLS_PSK_WITH_AES_256_CBC_SHA,
                      TLS_PSK_WITH_AES_256_CBC_SHA384,
TLS_PSK_WITH_AES_256_CCM,
                      TLS_PSK_WITH_AES_256_CCM_8,
                      TLS_PSK_WITH_AES_256_GCM_SHA384,
                      TLS_RSA_PSK_WITH_AES_128_CBC_SHA,
                      TLS_RSA_PSK_WITH_AES_128_CBC_SHA256,
                      TLS_RSA_PSK_WITH_AES_128_GCM_SHA256,
                      TLS_RSA_PSK_WITH_AES_256_CBC_SHA,
                      TLS_RSA_PSK_WITH_AES_256_CBC_SHA384,
                      TLS_RSA_PSK_WITH_AES_256_GCM_SHA384,
                      TLS_SRP_SHA_WITH_AES_128_CBC_SHA,
                      TLS_SRP_SHA_WITH_AES_256_CBC_SHA,
                      TLS_SRP_SHA_DSS_WITH_AES_128_CBC_SHA,
                      TLS_SRP_SHA_DSS_WITH_AES_256_CBC_SHA,
                      TLS_SRP_SHA_RSA_WITH_AES_128_CBC_SHA,
                      TLS_SRP_SHA_RSA_WITH_AES_256_CBC_SHA,
                      TLS_AES_128_GCM_SHA256, TLS_AES_256_GCM_SHA384,
                      TLS_CHACHA20_POLY1305_SHA256
```

Related Links

- security config modify

security config ocsp disable

Disable OCSP for one or more selected applications

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security config ocsp disable` command disables the OCSP-based certificate status check for applications supporting SSL/TLS communications. For more information about the OCSP-based certificate status check for applications supporting SSL/TLS communications, see the security config ocsp show command.

Parameters

`-application <SSL/TLS Application supporting OCSP>,…` - Application Name (privilege: advanced)

Use this parameter to specify the application to disable the OCSP support. To disable all applications, the value 'all' can be used. Note: You cannot specify the value 'all' with other applications.

Examples

The following example disables the OCSP support for AutoSupport and EMS applications:

```
cluster1::*> security config ocsp disable -application autosupport,ems

cluster1::> security config ocsp show
Application          OCSP Enabled?
-------------------- -------------
autosupport          false
audit_log            true
fabricpool           true
ems                  false
kmip                 true
ldap                 true
ssh                  true
6 entries were displayed.
```

The following example disables the OCSP support for all applications:

```
cluster1::*> security config ocsp disable -application all
Warning: OCSP will be disabled for all applications. Any previous
modifications
         will be ignored.
         Do you want to continue? {y|n}: y

cluster1::*> security config ocsp show
Application          OCSP Enabled?
-------------------- -------------
autosupport          false
audit_log            false
fabricpool           false
ems                  false
kmip                 false
ldap                 false
ssh                  false
6 entries were displayed.
```

Related Links

- security config ocsp show

security config ocsp enable

Enable OCSP for one or more selected applications

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security config ocsp enable` command enables the OCSP-based certificate status check for applications supporting SSL/TLS communications. For more information about the OCSP-based certificate status check for applications supporting SSL/TLS communications, see the security config ocsp show command.

Parameters

`-application <SSL/TLS Application supporting OCSP>,…` - List of Applications (privilege: advanced)

Use this parameter to specify the application to enable the OCSP support. To enable all applications, the value 'all' can be used. Note: You cannot specify the value 'all' with other applications.

Examples

The following example enables the OCSP support for AutoSupport and EMS applications:

```
cluster1::*> security config ocsp enable -application autosupport,ems

cluster1::> security config ocsp show
Application          OCSP Enabled?
-------------------- -------------
autosupport          true
audit_log            false
fabricpool           false
ems                  true
kmip                 false
ldap                 false
ssh                  true
6 entries were displayed.
```

The following example enables the OCSP support for all applications:

```
cluster1::*> security config ocsp enable -application all
Warning: OCSP will be enabled for all applications. Any previous
modifications
         will be ignored.
         Do you want to continue? {y|n}: y

cluster1::*> security config ocsp show
Application          OCSP Enabled?
-------------------- -------------
autosupport          true
audit_log            true
fabricpool           true
ems                  true
kmip                 true
ldap                 true
ssh                  true
6 entries were displayed.
```

Related Links

- security config ocsp show

security config ocsp show

Show Online Certificate Status Protocol (OCSP) settings

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security config ocsp show` command displays the support status of the OCSP-based certificate status check for applications supporting SSL/TLS communications. If the OCSP support is enabled for an application, this check is done in addition to the certificate chain validation as part of the SSL handshake process. The OCSP-based certificate status check is done for all the certificates in the chain, provided that each certificate has the OCSP URI access points mentioned in them. If no access points are specified, the OCSP-based certificate revocation status check is ignored for that certificate and checking continues for the rest of the certificates in the chain.

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.

[-application <SSL/TLS Application supporting OCSP>] - Application Name (privilege: advanced)

Selects the application that matches this parameter value. Applications include:

- autosupport - AutoSupport
- audit_log - Audit Logging
- fabricpool - External capacity tiers
- ems - Event Management System
- kmip - Key Management Interoperability Protocol
- ldap_ad - Lightweight Directory Access Protocol - Active Directory (query and modify items in Active Directory)
- ldap_nis_namemap - Lightweight Directory Access Protocol - NIS and Name Mapping (query Unix user, group, netgroup and name mapping information)
- ssh - Secure Shell

[-is-ocsp-enabled {true|false}] - Is OCSP-based Certificate Status Check Enabled? (privilege: advanced)

Selects the application that matches this parameter value.

Examples

The following example displays the OCSP support for the applications supporting SSL/TLS communications:

```
cluster1::> security config ocsp show
Application          OCSP Enabled?
-------------------- -------------
autosupport          true
audit_log            false
fabricpool           false
ems                  true
kmip                 false
ldap                 false
ssh                  false
6 entries were displayed.
```

The following example displays the OCSP support for AutoSupport:

```
cluster1::*> security config ocsp show -application autosupport
Application Name: autosupport
Is OCSP-based Certificate Status Check Enabled?: true
```

security config status show

(DEPRECATED)-Display Security Configuration Status

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

> This command has been deprecated as of 9.9 and may be removed in a future release of ONTAP. Reboot is no longer required to apply the security configuration, so it now always displays false.

```
The `security config status show` command displays the required reboot
status of the nodes in the cluster after security configuration settings
have been modified using the xref:{relative_path}security-config-
modify.html[security config modify] command. Use this command to monitor
the status of the required reboot process. When all nodes have rebooted,
the cluster is ready to use the new security configuration settings.
```

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>|local}] - Node Name (privilege: advanced)

Select the node whose reboot-status you want to display.

[-reboot-needed {true|false}] - Reboot Needed (privilege: advanced)

reboot-needed status of the node that tells if the node requires a reboot for security configuration to take effect.

Examples

The following example displays the status of a configuration change in a four-node cluster.

```
cluster1::> security config status show
Nodes in Cluster      Reboot Needed
--------------------- -------------------
node1                 true
node2                 true
node3                 false
node4                 false
4 entries were displayed.
```

The following example shows the output of the command after the cluster reboot process is complete.

```
cluster1::> security config status show
Nodes in Cluster      Reboot Needed
--------------------- -------------------
node1                 false
node2                 false
node3                 false
node4                 false
4 entries were displayed.
```

Related Links

- security config modify

security cryptomod-fips commands

security cryptomod-fips show

Display the status of cryptomod-fips

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays information about the status of the cryptomod FIPS module. By default, this command displays the following information:

- Node name
- FIPS version
- Module version
- FIPS state
- Boolean indicating if module is a user-space module
- Boolean indicating if module is operating in FIPS mode
- Boolean indicating if module is currently under validation

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>|local}] - Node

Selects the nodes that match this parameter value.

[-fips-state-text <text>] - FIPS State

Shows the FIPS state of the module.

- POWER ON STATE
- SELF-TEST STATE
- ERROR STATE
- OPERATE STATE
- POWER OFF STATE

[-fips-version <text>] - FIPS Version

Shows the FIPS version.

[-module-version <text>] - Module Version

Shows the cryptomod FIPS module version.

[-is-user-space-module {true|false}] - Is User Space Module?

True if the module is a user-space module.

[-is-fips-enabled {true|false}] - Is FIPS Mode Enabled?

True if the module is operating in FIPS mode.

[-is-iut-enabled {true|false}] - Is an IUT Module Enabled?

True if the module is currently under validation.

Examples

```
cluster1::> security cryptomod-fips show
Node   FIPS            FIPS    Module
       State           Version Version
------ --------------- ------- ---------
node-1 OPERATE STATE   140-2   2.2
node-2 OPERATE STATE   140-2   2.2
```

security dynamic-authorization commands

security dynamic-authorization modify

Modify dynamic-authorization global settings

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization modify` command modifies one or more dynamic authorization settings.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the setting. If this parameter is specified, the setting applies to that Vserver only. If not specified, the cluster Vserver is used.

`[-state {disabled|visibility|enforced}]` - Dynamic Authorization State

This parameter sets the state of the dynamic authorization feature. Valid values are *`disabled`*, *`visibility`* and *`enforced`*.

- disabled: Dynamic Authorization is disabled. This is the default factory setting.
- visibility: Dynamic Authorization is enabled in visibility mode. Customers will typically use this mode during a trial run to test the feature and ensure that users are not being inadvertently locked out. In this mode, the trust score is checked every time the user attempts to execute a restricted command, but not enforced. That is, the user will be allowed to execute all restricted commands as long as his RBAC privileges allow it. However, all commands that will either be denied or subject to additional MFA challenge will be logged.
- enforced: Dynamic Authorization is enabled in enforcement mode. Customers will typically use this mode after they have completed their trial run using visibility mode and verified that their configuration settings are correct, i.e. no users are being inadvertently locked out as a result of incorrect configuration. In this mode, the trust score is checked every time the user attempts to execute a restricted command and use to enforce dynamic authorization. That is, the user will be allowed to execute all restricted commands without additional MFA challenge only if the trust score exceeds the upper MFA challenge boundary. If the trust score falls within the lower and upper MFA challenge boundary, the user will be subject to an additional MFA challenge before being allowed to execute the command. If the trust score falls below the lower MFA challenge boundary, the user will be denied access. All additional MFA challenges and denials will be logged. The suppression interval is also enforced so no additional authentication challenges will be required if repeated authorization requests are made within the suppression interval.

`[-suppression-interval {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W | disabled}]` - Dynamic Authorization Suppression Interval

The dynamic authorization challenge suppression interval in ISO-8601 format. When a series of restricted commands are executed within a short interval, multiple authentication prompts are suppressed to create a good user experience. The default suppression interval is 10 minutes, or *`PT10M`* in ISO-8601 format.

`[-lower-challenge-boundary <percent>]` - Lower MFA Challenge Boundary

The lower MFA challenge percentage boundary. Supported values are from *`0`* to *`99`*. Default value is *`0`*.

`[-upper-challenge-boundary <percent>]` - Upper MFA Challenge Boundary

The upper MFA challenge percentage boundary. Supported values are from *`0`* to *`100`*. This must be equal to or greater than the value of the lower boundary. A value of 100 means that every request will either be denied or subject to an additional authentication challenge; there are no requests that are allowed without a challenge. Default value is *`90`*.

Examples

The following command modifies the lower challenge boundary to 10.

```
cluster1::> security dynamic-authorization modify -lower-challenge
-boundary 10

cluster1::> security dynamic-authorization show
Vserver: cluster1
                          Dynamic Authorization State: disabled
           Dynamic Authorization Suppression Interval: 10m
                        Lower MFA Challenge Boundary: 10%
                        Upper MFA Challenge Boundary: 90%
```

security dynamic-authorization show

Show dynamic-authorization global settings

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization show` displays information on dynamic authorization settings.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that you specify.

| [-instance] }

Displays all the fields for the dynamic authorization settings.

[-vserver <vserver name>] - Vserver

Selects the dynamic authorization settings that match this parameter value. If not specified, all cluster-level and Vserver-level settings are displayed.

[-state {disabled|visibility|enforced}] - Dynamic Authorization State

Selects the dynamic authorization settings that match this parameter value.

[-suppression-interval {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W | disabled}] - Dynamic Authorization Suppression Interval

Selects the dynamic authorization settings that match this parameter value.

[-lower-challenge-boundary <percent>] - Lower MFA Challenge Boundary

Selects the dynamic authorization settings that match this parameter value.

[-upper-challenge-boundary <percent>] - Upper MFA Challenge Boundary

Selects the dynamic authorization settings that match this parameter value.

Examples

The example below displays information on dynamic authorization settings:

```
cluster1::> security dynamic-authorization show
Vserver: cluster1
                           Dynamic Authorization State: disabled
            Dynamic Authorization Suppression Interval: 10m
                         Lower MFA Challenge Boundary: 0%
                         Upper MFA Challenge Boundary: 90%
```

security dynamic-authorization authentication-history-policy modify

Modify authentication history policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization authentication-history-policy modify` command updates the authentication history policy settings for dynamic authorization.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the authentication history policy setting. If this parameter is specified, the setting applies to that Vserver only. If not specified, the cluster Vserver setting is used.

`[-use-last-num-days <integer>]` - Last Number of Days

This parameter optionally specifies the last number of days of authentication history statistics to use in calculating the trust score for the authentication history component. By default, this is set to -1, which means the trust score for authentication history component is calculated from all successful and failed authentications since the user's first successful login.

`[-lower-boundary <percent>]` - Lower Boundary of Authentication Failures

This parameter optionally specifies the lower boundary of authentication failures. The value is a percentage from 0 to 99, and must be less than or equal to the upper boundary. When used in conjunction with the *`upper-boundary`* , if the authentication failures are less than the *`lower-boundary`* percentage, the authentication history component gets a full trust score, while if the authentication failures are higher than the *`upper-boundary`* percentage, the authentication history component gets a zero trust score. Authentication failures falling between the *`lower-boundary`* and *`upper-boundary`* gets a 50% trust score for the authentication history component.

`[-upper-boundary <percent>]` - Upper Boundary of Authentication Failures

This parameter optionally specifies upper boundary of authentication failures. The value is a percentage from 0 to 100, and must be greater than or equal to the lower boundary. Refer to the description in the *`lower-boundary`* parameter on how this setting is used.

Examples

The following command modifies the upper boundary of authentication failures for the Administrative Vserver to 90%.

```
cluster1::*> security dynamic-authorization authentication-history-policy
modify -upper-boundary 90

cluster1::*> security dynamic-authorization authentication-history-policy
show
Vserver: cluster1
                                  Last Number of Days: 90
            Lower Boundary of Authentication Failures: 10%
            Upper Boundary of Authentication Failures: 90%
Vserver: svm0
                                  Last Number of Days: -1
            Lower Boundary of Authentication Failures: 10%
            Upper Boundary of Authentication Failures: 100%
2 entries were displayed.
```

security dynamic-authorization authentication-history-policy show

Show authentication history policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization authentication-history-policy show` displays information about the dynamic authorization authentication history policy settings.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that you specify.

| [-instance] }

Displays all the fields for the dynamic authorization authentication history policy.

[-vserver <vserver name>] - Vserver

Selects the dynamic authorization authentication history policy fields that match this parameter value.

[-use-last-num-days <integer>] - Last Number of Days

Selects the dynamic authorization authentication history policy fields that match this parameter value.

[-lower-boundary <percent>] - Lower Boundary of Authentication Failures

Selects the dynamic authorization authentication history policy fields that match this parameter value.

[-upper-boundary <percent>] - Upper Boundary of Authentication Failures

Selects the dynamic authorization authentication history policy fields that match this parameter value.

Examples

The example below displays information about all dynamic authorization authentication history policy settings:

```
cluster1::> security dynamic-authorization authentication-history-policy
show
Vserver: cluster1
                          Last Number of Days: 90
    Lower Boundary of Authentication Failures: 10%
    Upper Boundary of Authentication Failures: 100%
Vserver: svm0
                          Last Number of Days: -1
    Lower Boundary of Authentication Failures: 10%
    Upper Boundary of Authentication Failures: 100%
2 entries were displayed.
```

security dynamic-authorization executed-commands show

Display executed commands

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization executed-commands show` command displays information about the executed commands according to the dynamic authorization rules.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that you specify.

| [-instance] }

Displays all the fields for the dynamic authorization executed commands.

[-vserver <vserver name>] - Vserver

Selects the dynamic authorization executed command fields that match this Vserver.

[-date <MM/DD/YYYY HH:MM:SS>] - Date

Selects the dynamic authorization executed command fields that match this date.

[-username <text>] - Username

Selects the dynamic authorization executed command fields that match this username.

[-operation <text>] - Operation

Selects the dynamic authorization executed command fields that match this operation.

[-count <integer>] - Count

Selects the dynamic authorization executed command fields that match this count.

[-score <integer>] - Trust Score

Selects the dynamic authorization executed command fields that match this score.

[-result {permit|deny|challenge}] - Result

Selects the dynamic authorization executed command fields that match this result.

Examples

The example below displays information about all dynamic authorization executed commands:

```
cluster1::> security dynamic-authorization executed-commands show

Vserver: usernamecluster-1

Date                  Operation   Username   Count   Trust Score   Result
-------------------   ---------   --------   -----   -----------   ------
12/7/2023 08:25:57    security login create
                                  admin      1       100           permit
12/7/2023 08:26:04    security login unlock
                                  admin      1       100           permit
12/7/2023 08:26:09    security multi-admin-verify approval-group create
                                  admin      1       100           permit
3 entries were displayed.
```

security dynamic-authorization group create

Add a Dynamic Authorization group

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization group create` command creates the groups to include in dynamic authorization.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver of the dynamic authorization group that is being created. If this parameter is specified, the setting applies to that Vserver only. If not specified, the cluster Vserver setting is used.

-name <text> - Group Name

This parameter specifies the name of the group that will be part of dynamic authorization.

[-excluded-usernames <text>,…] - List of Excluded Users

This parameter optionally specifies the list of users that will be excluded from dynamic authorization.

[-comment <text>] - Comment

This parameter optionally specifies the comments.

Examples

The following command creates a group *test* on vserver *vs1* and excludes the user *tsmith* from dynamic authorization.

```
cluster1::> security dynamic-authorization group create -vserver vs1 -name
test -excluded-usernames tsmith
```

security dynamic-authorization group delete

Delete a Dynamic Authorization group

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization group delete` command deletes the specified group.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver of the dynamic authorization group that is being deleted. If this parameter is specified, the setting applies to that Vserver only. If not specified, the cluster Vserver setting is used.

-name <text> - Group Name

This parameter specifies the group name that is being deleted.

Examples

The following command deletes the group *test* from the Vserver *vs1* .

```
cluster1::> dynamic authorization group delete -vserver vs1 -name test
```

security dynamic-authorization group modify

Modify a Dynamic Authorization group

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization group modify` command modifies the dynamic authorization groups.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver of the group for which the `-excluded-usernames` or `-comment` is being modified. If this parameter is specified, the setting applies to that Vserver only. If not specified, the cluster Vserver setting is used.

`-name <text>` - Group Name

This parameter specifies the name of the group for which the `-excluded-usernames` or `-comment` is being modified.

`[-excluded-usernames <text>,…]` - List of Excluded Users

This parameter specifies the list of users to be excluded from dynamic authorization.

`[-comment <text>]` - Comment

This parameter optionally specifies the comments.

Examples

The following command modifies the excluded users for the group *`test`* who is part of Vserver *`vs1`* .

```
cluster1::> security dynamic-authorization group modify -vserver vs1
-group-name test -excluded-usernames Jsmith
```

security dynamic-authorization group show

Display Dynamic Authorization groups

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization group show` command displays information about the dynamic authorization groups.

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.

[-vserver <vserver name>] - Vserver

Selects dynamic-authorization groups for this Vserver.

[-name <text>] - Group Name

Selects dynamic-authorization groups that match this group name.

[-excluded-usernames <text>,…] - List of Excluded Users

Selects the dynamic-authorization groups that match the specified excluded username.

[-comment <text>] - Comment

Selects the dynamic-authorization groups that match this comment.

Examples

The example below displays dynamic authorization group information for the Vserver *vs1* .

```
cluster1::> security dynamic-authorization group show -vserver vs1
               Vserver: vs1
            Group Name: NETAPP_ENG
List of Excluded Users: user1, user2, user12
               Comment: -
```

security dynamic-authorization rule create

Add a dynamic authorization rule

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization rule create` command creates a custom dynamic authorization rule for an operation. By default, the set of operations subject to dynamic authorization is the same as the default Multi-Admin-Verify (MAV) set of commands. Additional operations can be configured using the `security dynamic-authorization rule create` command.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver associated with the custom dynamic authorization rule.

-operation <Command or Command Directory> - Command or Command Directory

This parameter specifies the operation for the custom dynamic authorization rule to be created. The operation can be a command or command directory.

-query <query> - Query

This parameter optionally specifies the object (or objects) upon which to apply the operation. Any field or query supported by the operation can be supplied. If a query is not specified for the rule, the rule applies to all objects of the specified operation. The query object must be enclosed in double quotation marks ("").

Examples

The following command creates a custom dynamic authorization rule for the job delete operation for the Administrative Vserver. This rule is applicable only to job objects whose job ID is greater than 50.

```
cluster1::> security dynamic-authorization rule create -operation "job
delete" -query "-id >50"
```

The following command creates a custom dynamic authorization rule for the snapmirror policy create operation for the data Vserver *vs1.example.com* . This rule is applicable only to SnapMirror policies of type other than *async-mirror* .

```
cluster1::> security dynamic-authorization rule create -vserver
vs1.example.com -operation "snapmirror policy create" -query "-type
!async-mirror"
```

Related Links

- job delete
- snapmirror policy create

security dynamic-authorization rule delete

Delete a dynamic authorization rule

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization rule delete` command deletes a custom dynamic authorization rule for an operation. It can be used to delete a custom dynamic authorization rule that was configured using the security dynamic-authorization rule create command.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the custom dynamic authorization rule.

`-operation <Command or Command Directory>` - Command or Command Directory

This parameter specifies the operation for the custom dynamic authorization rule to be deleted.

Examples

The following command deletes a custom dynamic authorization rule for the `network port ifgrp` operation for the Administrative Vserver.

```
cluster1::> security dynamic-authorization rule delete -vserver cluster1
-operation "network port ifgrp"
```

The following command deletes a custom dynamic authorization rule for the `vserver services nis-domain create` operation for the data Vserver *`vs1.example.com`*.

```
cluster1::> security dynamic-authorization rule delete -vserver
vs1.example.com -operation "vserver services nis-domain create"
```

Related Links

- security dynamic-authorization rule create

security dynamic-authorization rule modify

Modify a dynamic authorization rule

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization rule modify` command modifies a custom dynamic authorization rule for an operation. It can be used to modify a custom dynamic authorization rule that was configured using the security dynamic-authorization rule create command.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the custom dynamic authorization rule.

`-operation <Command or Command Directory>` - Command or Command Directory

This parameter specifies the operation for the custom dynamic authorization rule to be modified. The operation can be a command or command directory.

`[-query <query>]` - Query

This parameter optionally specifies the object (or objects) upon which to apply the operation. Any field or query supported by the operation can be supplied. If the query is specified as "" i.e., empty, the rule applies to all objects of the specified operation. The query object must be enclosed in double quotation marks ("").

Examples

The following command modifies the query of a custom dynamic authorization rule for the storage encryption disk destroy operation in the Administrative Vserver. The new query disallows destroying of storage encryption disks starting with the name `xxxxx_` .

```
cluster1::> security dynamic-authorization rule modify -operation "storage
encryption disk destroy" -query "-disk !xxxxx_*"
```

The following command resets the query of a custom dynamic authorization rule for the vserver active-directory create operation for the data Vserver *vs1.example.com*.

```
cluster1::> security dynamic-authorization rule modify -vserver
vs1.example.com -operation "vserver active-directory create" -query ""
```

Related Links

- security dynamic-authorization rule create
- storage encryption disk destroy
- vserver active-directory create

security dynamic-authorization rule show

Show dynamic authorization rules

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization rule show` displays information about dynamic authorization rules, which includes both pre-defined as well as custom dynamic authorization rules.

Parameters

{ [-fields <fieldname>,…]
Selects the fields that you specify.

| [-instance] }
Displays all the fields for the specified dynamic authorization rules.

[-vserver <vserver name>] - Vserver
Selects the dynamic authorization rules that match this parameter value.

[-operation <Command or Command Directory>] - Command or Command Directory
Selects the dynamic authorization rules that match this parameter value.

[-query <query>] - Query
Selects the dynamic authorization rules that match this parameter value.

Examples

The example below displays information about all dynamic authorization rules whose operation matches the prefix *security multi-admin-verify*.

```
cluster1::*> security dynamic-authorization rule show -operation "security
multi-admin-verify"*

Vserver: cluster1

Operation
Query
----------------------------------------
-----------------------------------
security multi-admin-verify
security multi-admin-verify approval-group
security multi-admin-verify approval-group replace
security multi-admin-verify rule

Vserver: vs1

Operation
Query
----------------------------------------
-----------------------------------
security multi-admin-verify
security multi-admin-verify approval-group
security multi-admin-verify approval-group replace
security multi-admin-verify rule
8 entries were displayed.
```

security dynamic-authorization trust-score-component create

Create a trust score component

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization trust-score-component create` command creates and registers a custom trust score component. Administrators can use this command to configure trust score components in addition to or as an alternative to built-in components.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the custom trust score component. If this parameter is specified, the setting applies to that Vserver only. If not specified, the cluster Vserver setting is used.

`-component <text>` - Component Name

The name of the custom component used to obtain the trust score. This must be unique within the Vserver.

[-weight <integer>] - Score Weight

An integer giving the raw weight of the component, indicating the importance of the component relative to other components for calculating the trust score. Built-in components have a default weightage of *20* .

[-provider-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}] - Trust Score Provider URI of Component

The trust score provider URI to obtain the trust score for the component. The response from the URI must be in JSON.

[-max-score <integer>] - Max Trust Score of Component

The maximum score for the component. The default value is *20* .

[-min-score <integer>] - Min Trust Score of Component

The minimum score for the component. The default value is *0* .

[-score-field <text>] - Score field to check in JSON response

The field within the JSON response to obtain the trust score.

[-score-type {trust-score|risk-score}] - Score Type

This parameter specifies if the score returned from the component is trust score or risk score. The trust score is in ascending order with a higher score denoting a higher trust level, while the risk score is in descending order. The default value is *trust-score* .

[-secret-access-key <text>] - Access key for trust score provider

An optional field giving the access key for the trust score provider. This is used to authenticate to the provider.

[-provider-http-headers <text>,…] - Provider HTTP headers

An optional list of HTTP headers required by the trust score provider.

Examples

The following command creates a dynamic authorization custom component for the Administrative Vserver. The username is a parameter that will be replaced with the actual username at run-time:

```
cluster1::> security dynamic-authorization trust-score-component create
-component comp1 -weight 20 -max-score 500 -provider-uri
https://provider.example.com/trust-scores/users/${username}/component
-score-field score
```

security dynamic-authorization trust-score-component delete

Delete a trust score component

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization trust-score-component delete` command deletes a custom trust score component. It can be used to delete a custom trust score component that was configured using the security dynamic-authorization trust-score-component create command.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the custom trust score component. If not specified, this defaults to the cluster Vserver.

`-component <text>` - Component Name

Name of the custom trust score component to be deleted.

Examples

The following command deletes a custom trust score component named comp1 for the Administrative Vserver.

```
cluster1::> security dynamic-authorization trust-score-component delete
-component comp1
```

Related Links

- security dynamic-authorization trust-score-component create

security dynamic-authorization trust-score-component modify

Modify a trust score component

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization trust-score-component modify` command updates an existing custom trust score component.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the custom trust score provider component. If this parameter is specified, the setting applies to that Vserver only. If not specified, the cluster Vserver setting is used.

`-component <text>` - Component Name

The component name.

`[-weight <integer>]` - Score Weight

An integer giving the raw weight of the component, indicating the importance of the component relative to other components for calculating the trust score. Built-in components have a default weightage of 20.

[-provider-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}] - Trust Score Provider URI of Component

The trust score provider URI to obtain the trust score for the component. The response from the URI must be in JSON.

[-max-score <integer>] - Max Trust Score of Component

The maximum score for the component.

[-min-score <integer>] - Min Trust Score of Component

The minimum score for the component.

[-score-field <text>] - Score field to check in JSON response

The field within the JSON response to obtain the trust score.

[-score-type {trust-score|risk-score}] - Score Type

This parameter specifies if the score returned from the component is trust score or risk score. The trust score is in ascending order with a higher score denoting a higher trust level, while the risk score is in descending order. The default value is *trust-score* .

[-secret-access-key <text>] - Access key for trust score provider

An optional field giving the access key for the trust score provider. This is used to authenticate to the provider.

[-provider-http-headers <text>,…] - Provider HTTP headers

An optional list of HTTP headers required by the trust score provider.

Examples

The following command modifies a dynamic authorization custom component for the Administrative Vserver to change the weightage of the component to 100.

```
cluster1::> security dynamic-authorization trust-score-component modify
-component comp1 -weight 100
```

security dynamic-authorization trust-score-component show

Display trust score components

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization trust-score-component show` displays information about the components of the trust score.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that you specify.

| [-instance] }

Displays all the fields for the specified dynamic authorization trust score components.

[-vserver <vserver name>] - Vserver

Selects the dynamic authorization trust score components that match this parameter value.

[-component <text>] - Component Name

Selects the dynamic authorization trust score components that match this parameter value.

[-weight <integer>] - Score Weight

Selects the dynamic authorization trust score components that match this parameter value.

[-max-percent-score-weight <double hundredths>] - Max Percentage Score Weight

Selects the dynamic authorization trust score components that match this parameter value.

[-provider-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}] - Trust Score Provider URI of Component

Selects the dynamic authorization trust score components that match this parameter value.

[-max-score <integer>] - Max Trust Score of Component

Selects the dynamic authorization trust score components that match this parameter value.

[-min-score <integer>] - Min Trust Score of Component

Selects the dynamic authorization trust score components that match this parameter value.

[-score-field <text>] - Score field to check in JSON response

Selects the dynamic authorization trust score components that match this parameter value.

[-score-type {trust-score|risk-score}] - Score Type

This parameter specifies if the score returned from the components is trust score or risk score. The trust score is in ascending order with a higher score denoting a higher trust level, while the risk score is in descending order. The default value is *trust-score* .

[-provider-http-headers <text>,…] - Provider HTTP headers

Selects the dynamic authorization trust score components that match this parameter value.

Examples

The example below displays information about all dynamic authorization trust score components, both built-in and custom:

```
cluster1::> security dynamic-authorization trust-score-component create
-component comp1 -weight 20 -max-score 100 -provider-uri
https://provider.example.com/trust-scores/users/admin1/component1.json
-score-field score

cluster1::> security dynamic-authorization trust-score-component show
Percentage
Vserver           Component Name                   Score Weight   Score
Weight
---------------   ------------------------------   ------------
------------
cluster1          authentication_history_policy              20
33.33
cluster1          comp1                                      20
33.33
cluster1          trusted_device                             20
33.33
svm0              authentication_history_policy              20
50.00
svm0              trusted_device                             20
50.00

5 entries were displayed.
```

The following command displays the details of all components matching the name *comp1* :

```
cluster1::> security dynamic-authorization trust-score-component show
-vserver cluster1 -component comp1 -instance
Vserver: cluster1
            Trust Score Component Name: comp1
               Weight of the Component: 20
Max Percentage Weight of the component: 50.00
 Trust Score Provider URI of Component:
https://provider.example.com/trust-scores/users/admin1/component1.json
                Max Score of Component: 100
 Score field to check in JSON response: score
                 Provider HTTP headers: -
```

security dynamic-authorization user-trust-score reset

Resets trust score of user

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security dynamic-authorization user-trust-score reset` command resets the trust score of the specified user.

Parameters

`-vserver <vserver name>` - Vserver

Selects the Vserver that match this parameter value.

`-username <text>` - Username

Reset the trust score for this user.

`-component <text>` - Component Name

The component for which the user trust score has to be reset.

Examples

The example below resets the user trust score.

```
cluster1::> security dynamic-authorization user-trust-score reset -vserver
vs1 -username Tsmith -component authentication_history_policy
```

security ipsec commands

security ipsec show-ikesa

Show IKE SA Information

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security ipsec show-ikesa` command displays information about IKE Security Associations (SA).

Running this command with the `-node` parameter displays information relevant to IKE SAs generated at the specified node.

Running this command with the `-vserver` parameter displays information relevant to IKE SAs associated with the specified vserver.

Running this command with the `-policy-name` parameter displays information relevant to IKE SAs created based on the specified security policy.

You can specify additional parameters to display only information matching those parameters. For example, to display IKE SAs associated with a specific local address, run the command with the `-local-address` parameter.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the specified fields. Notice that key fields are always displayed.

| [-instance] }

If you specify the `-instance` parameter, the command displays all fields of the IKE SAs.

-node <nodename> - Node

This required parameter specifies the node from which the IKE SA information will be collected and displayed.

[-vserver <vserver name>] - Vserver Name

Use this parameter to display the IKE SAs associated with the specified Vserver.

[-policy-name <text>] - Policy Name

Use this parameter to display the IKE SAs created based on the specified security policy.

[-local-address <text>] - Local Address

Use this parameter to display the IKE SAs with the specified local endpoint IP address.

[-remote-address <text>] - Remote Address

Use this parameter to display the IKE SAs with the specified remote endpoint IP address.

[-initiator-spi <text>] - Initiator SPI

Use this parameter to display the IKE SAs with the specified initiator Security Parameter Index (SPI).

[-responder-spi <text>] - Responder SPI

Use this parameter to display the IKE SAs with the specified responder SPI.

[-is-initiator {true|false}] - Is Initiator

Use this parameter to display the IKE SAs created when the given node matches the specified initiator role: true means initiator role and false means responder role in IKE negotiation.

[-ike-version <integer>] - IKE Version

Use this parameter to display the IKE SAs created using the specified IKE version.

[-auth-method <IKE Authentication Method>] - Authentication Method

Use this parameter to display the IKE SAs created using the specified authentication method.

[-state <IKE SA State>] - IKE SA State

Use this parameter to display only the IKE SAs that are in the specified state.

[-cipher-suite <Cipher Suite Type>] - Cipher Suite

Use this parameter to display the IKE SAs created using the specified cipher suite.

[-lifetime <integer>] - Lifetime

Use this parameter to display the IKE SAs with the specified remaining lifetime. Notice that lifetime keeps changing for the duration of the security association.

Examples

This example displays all IKE SAs for node *cluster1-node1* :

```
cluster-1::> security ipsec show-ikesa -node cluster1-node1
            Policy Local           Remote
Vserver     Name   Address         Address         Initator-SPI     State
----------- ------ --------------- --------------- ----------------
-----------
vs1         Policy1
                   192.186.10.1    192.186.10.2    e658e5bc7ece199e
ESTABLISHED
vs2         Policy2
                   192.168.20.1    192.168.20.2    8eac392028ab4f12
ESTABLISHED
2 entries were displayed.
```

This example displays selected fields of all IKE SAs for node *cluster1-node1* :

```
cluster-1::> security ipsec show-ikesa -node cluster1-node1 -fields is-
initiator,initiator-spi,responder-spi,auth-method,cipher-suite,lifetime

node           vserver policy-name local-address remote-address initiator-
spi    responder-spi    is-initiator auth-method cipher-suite  lifetime
-------------- ------- ----------- ------------- --------------
---------------- ---------------- ------------ ----------- -------------
--------
cluster1-node1 vs1     Policy1     192.186.10.1  192.186.10.2
e658e5bc7ece199e 9b61befff71e8ca2 false        PSK         SUITEB_GCM256
6300
cluster1-node1 vs2     Policy2     192.186.20.1  192.186.20.2
4d43aaba8ca01cd8 00bdd5aac569e08a true         PSK         SUITEB_GCM256
6720
2 entries were displayed.
```

This example displays all IKE SAs for vserver *vs1* :

```
cluster-1::> security ipsec show-ikesa -node cluster1-node1
            Policy Local           Remote
Vserver     Name   Address         Address         Initator-SPI     State
----------- ------ --------------- --------------- ----------------
-----------
vs1         Policy1
                   192.186.10.1    192.186.10.2    e658e5bc7ece199e
ESTABLISHED
```

This example displays instance view (all fields) for all IKE SAs associated with node *cluster1-node1* , vserver *vs1* and created using policy *Policy1* :

```
cluster-1::> security ipsec show-ikesa -node cluster1-node1 -vserver vs1
-policy-name Policy1 -instance
Node: cluster1-node1
         Vserver Name: vs1
          Policy Name: Policy1
        Local Address: 192.168.10.1
       Remote Address: 192.168.10.2
        Initiator SPI: e658e5bc7ece199e
        Responder SPI: 9b61befff71e8ca2
         Is Initiator: false
          IKE Version: 2
Authentication Method: PSK
         IKE SA State: ESTABLISHED
         Cipher Suite: SUITEB_GCM256
             Lifetime: 6000
```

security ipsec show-ipsecsa

Show IPsec SA Information

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security ipsec show-ipsecsa` command displays information about IPsec Security Associations (SA).

Running the command with the `-node` parameter displays information relevant to IPsec SAs at the specified node.

Running this command with the `-vserver` parameter displays information relevant to IPsec SAs associated with the specified vserver.

Running this command with the `-policy-name` parameter displays information relevant to IPsec SAs created

using the specified security policy.

You can specify additional parameters to display only information matching those parameters. For example, to display IPsec SAs only about a certain local address, run the command with the `-local-address` parameter.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the specified fields. Notice that key fields are always displayed.

| [-instance] }

If you specify the `-instance` parameter, the command displays all fields of matching IPsec SAs.

-node <nodename> - Node

This required parameter specifies from which node the IPsec SA information will be collected and displayed.

[-vserver <vserver name>] - Vserver Name

Use this parameter to display the IPsec SAs associated with the specified Vserver.

[-policy-name <text>] - Policy Name

Use this parameter to display the IPsec SAs created based on the specified security policy.

[-local-address <text>] - Local Address

Use this parameter to display the IPsec SAs with the specified local endpoint IP address.

[-remote-address <text>] - Remote Address

Use this parameter to display the IPsec SAs with the specified remote endpoint IP address.

[-inbound-spi <text>] - Inbound SPI

Use this parameter to display the IPsec SA having the specified inbound Security Parameter Index (SPI).

[-outbound-spi <text>] - Outbound SPI

Use this parameter to display the IPsec SA having the specified outbound SPI.

[-action <IPsec Action Type>] - IPsec Action

Use this parameter to display IPsec SAs with the specified security action type, such as ESP_TRA for ESP transport mode protection or BYPASS to bypass IPsec, or DISCARD.

[-state <text>] - IPsec SA State

Use the parameter to display only the IPsec SAs that are in the specified state.

[-cipher-suite <Cipher Suite Type>] - Cipher Suite

Use this parameter to display the IPsec SAs that use the specified cipher-suite.

[-ib-bytes <integer>] - Inbound Bytes Processed

Use this parameter to display the IPsec SAs matching the processed inbound bytes. Notice that ib-bytes keeps changing as inbound packets are processed.

[-ib-pkts <integer>] - Inbound Pkts Processed

Use this parameter to display the IPsec SAs matching the processed inbound packets. Notice that ib-pkts keeps changing as inbound packets are processed.

[-ob-bytes <integer>] - Outbound Bytes Processed

Use this parameter to display the IPsec SAs matching the processed outbound bytes. Notice that ob-bytes keeps changing as outbound packets are processed.

[-ob-pkts <integer>] - Outbound Pkts Processed

Use this parameter to display the IPsec SAs matching the processed outbound packets. Notice that ob-pkts keeps changing as outbound packets are processed.

[-lifetime <integer>] - IPsec SA Lifetime Seconds

Use this parameter to display the IPsec SAs matching the remaining lifetime. Notice that lifetime keeps changing for the duration of the security association.

[-hw-ib-bytes <integer>] - Inbound Bytes Processed by Offload

Use this parameter to display the IPsec SAs matching the processed inbound bytes by hardware offload. Notice that number of bytes keeps changing as inbound packets are processed.

[-hw-ib-pkts <integer>] - Inbound Pkts Processed by Offload

Use this parameter to display the IPsec SAs matching the processed inbound packets by hardware offload. Notice that number of packets keeps changing as inbound packets are processed.

[-hw-ob-bytes <integer>] - Outbound Bytes Processed by Offload

Use this parameter to display the IPsec SAs matching the processed outbound bytes by hardware offload. Notice that number of bytes keeps changing as outbound packets are processed.

[-hw-ob-pkts <integer>] - Outbound Pkts Processed by Offload

Use this parameter to display the IPsec SAs matching the processed outbound packets by hardware offload. Notice that number of packets keeps changing as outbound packets are processed.

Examples

The this example displays all IPsec SAs for node *cluster1-node1* :

```
cluster-1::> security ipsec show-ipsecsa -node cluster1-node1
            Policy  Local           Remote          Inbound  Outbound
Vserver     Name    Address         Address         SPI      SPI
State
----------- ------- --------------- --------------- -------- --------
---------
vs1         Policy1
                    192.186.10.1    192.186.10.2    c68de9db c84f913b
INSTALLED
vs2         Policy2
                    192.186.20.1    192.186.20.2    cbc01493 c6ee7424
INSTALLED
2 entries were displayed.
```

This example displays selected fields of all IPsec SAs for node *cluster1-node1* :

```
cluster-1::> security ipsec show-ipsecsa -node cluster1-node1 -fields
local-address,remote-address,inbound-spi,outbound-spi
node           vserver policy-name local-address  remote-address inbound-
spi outbound-spi
-------------- ------- ----------- -------------- --------------
----------- ------------
cluster1-node1 vs1     Policy1     192.186.10.1   192.186.10.2   c68de9db
c84f913b
cluster1-node1 vs2     Policy2     192.186.20.1   192.186.20.2   cbc01493
c6ee7424
2 entries were displayed.
```

```
This example displays selected fields of all IPsec SAs associated with
node ``_cluster1-node1_``:
cluster-1::> security ipsec show-ipsecsa -node cluster1-node1 -fields ib-
bytes,ib-pkts,ob-bytes,ob-pkts
node           vserver policy-name local-address  remote-address inbound-
spi ib-bytes ib-pkts ob-bytes ob-pkts
-------------- ------- ----------- -------------- --------------
----------- -------- ------- -------- -------
cluster1-node1 vs1     Policy1     192.186.10.1   192.186.10.2   c68de9db
4704     56      6720     56
cluster1-node1 vs2     Policy2     192.186.20.1   192.186.20.2   cbc01493
20434    115     23082    120
2 entries were displayed.
```

This example displays instance view (all fields) for all IPsec SAs associated with node *cluster1-node1* ,

vserver *vs1* and created using policy *Policy1* :

```
cluster-1::> security ipsec show-ipsecsa -node cluster1-node1 -vserver vs1
-policy-name Policy1 -instance
Node: cluster1-node1
             Vserver Name: vs1
              Policy Name: Policy1
              Inbound SPI: c68de9db
             Outbound SPI: c84f913b
            Local Address: 192.168.10.1
           Remote Address: 192.168.10.2
             IPsec Action: ESP_TRA
           IPsec SA State: INSTALLED
             Cipher Suite: SUITEB_GCM256
  Inbound Bytes Processed: 4704
   Inbound Pkts Processed: 56
 Outbound Bytes Processed: 6720
  Outbound Pkts Processed: 56
IPsec SA Lifetime Seconds: 1800
```

security ipsec ca-certificate add

Add CA certificate(s) to a vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command adds a list of CA certificates to IPsec for the given Vserver. These certificates will be used for PKI authentication with remote IKE endpoint. The CA certificates should have already been installed using either security certificate install command or security certificate create command.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver for which the IPsec CA certificates should be added.

`-ca-certs <text>,…` - CA Certificate Names

Use this parameter to specify the list of CA certificates to be added to IPsec.

Examples

The following example adds two IPsec CA certificates named caCert1 and caCert2 to Vserver v1.

```
cluster-1::>security ipsec ca-certificate add -vserver v1 -ca-certs
caCert1,caCert2
```

Related Links

- security certificate install
- security certificate create

security ipsec ca-certificate remove

Remove CA certificate(s) from a vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command removes a list of IPsec CA certificates for the given Vserver. The CA certificates being removed should have been previously added to IPsec using security ipsec ca-certificate add command.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver for which the IPsec CA certificates should be removed.

`-ca-certs <text>,…` - CA Certificate Names

Use this parameter to specify the list of CA certificates to be removed from IPsec.

Examples

The following example removes two IPsec CA certificates named caCert1 and caCert2 for Vserver v1.

```
cluster-1::>security ipsec ca-certificate remove -vserver v1 -ca-certs
caCert1,caCert2
```

Related Links

- security ipsec ca-certificate add

security ipsec ca-certificate show

Displays the CA certificates added to IPsec

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the configured IPsec CA certificates.

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.

[-vserver <vserver name>] - Vserver Name

If you specify this parameter, then the command displays only the IPsec CA certificates configured for the given vserver.

[-ca-certs <text>,…] - CA Certificate Names

If you specify this parameter, then the command displays only the Vservers for which the given CA certificates are present in IPsec.

Examples

The following example lists the IPsec CA certificates configured for all Vservers.

```
cluster-1::>security ipsec ca-certificate show

Vserver              CA Certificate Names
------------------   ---------------------------------
v1                   caCert1, caCert2
v2                   caCert3, caCert4
2 entries were displayed.
```

security ipsec config modify

Modify IPsec config

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command modifies IPsec configuration parameters.

Parameters

[-is-enabled {true|false}] - Is IPsec Enabled

This parameter enables and disables IPsec on the storage system.

[-log-level <IPsec Log Level>] - IPsec Logging Level

This parameter sets the IPsec logging level, where logging level 0 means no logging, and logging level 5 is most verbose. Default value is 2.

[-replay-window {0|64|128|256|512|1024}] - IPsec Replay Window Size

This parameter sets the IPsec replay window size. The possible values are 0, 64, 128, 256, 512 and 1024. Default value is 0.

[-ready-to-downgrade {true|false}] - IPsec Ready To Downgrade (privilege: advanced)

This parameter is used when downgrading to a non-IPsec capable ONTAP. Set this parameter to true to cleanup IPsec configurations before such downgrade.

`[-is-offload-enabled {true|false}]` - Offload Enabled

This parameter enables and disables IPsec NIC offload. If enabled, ONTAP will try to use NIC offload whenever possible.

Examples

The following example enables IPsec:

```
cluster-1::> security ipsec config modify -is-enabled true
```

The following example sets the IPsec logging level to 4:

```
cluster-1::> security ipsec config modify -log-level 4
```

The following example sets the IPsec replay window size to 64:

```
cluster-1::> security ipsec config modify -replay-window 64
```

The following example enables the IPsec NIC offload. IPsec itself has to be enabled to use the IPsec NIC offload:

```
cluster-1::> security ipsec config modify -is-enabled true -is-offload
-enabled true
```

security ipsec config show

Display IPsec config

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command shows the current IPsec configuration parameters.

Examples

The following example shows the state of IPsec (enabled/disabled) and the IPsec logging level:

```
cluster-1::> security ipsec config show
     IPsec Enabled: false
   IPsec Log Level: 2
Replay Window Size: 0
   Offload Enabled: true
```

security ipsec policy create

Create an IPsec policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command creates a new IPsec policy. The traffic to be protected is specified by the traffic selectors (local-ip-subnets, remote-ip-subnets, local-ports, remote-ports, protocols). IPsec is not supported for the admin Vserver in a MetroCluster environment.

Parameters

`-vserver <vserver name>` - Vserver

Specifies the Vserver to which the policy will belong. If there is only a single Vserver capable of supporting IPsec, the Vserver parameter is implied.

`-name <text>` - Policy Name

This required parameter specifies the name of the policy which may be a text string (1-64 bytes), a hexadecimal string (beginning with '0x') or a base64 encoded binary string (beginning with '0s').

`-local-ip-subnets <IP Address/Mask>,…` - Local IP Subnets

This required parameter specifies the IPv4 or IPv6 subnet (address and mask, can be subnet or individual address) representing the local address (range) to be protected by this policy.

`-remote-ip-subnets <IP Address/Mask>,…` - Remote IP Subnets

This required parameter specifies the IPv4 or IPv6 subnet (address and mask, can be subnet or individual address) representing the remote address (range) to be protected by this policy.

`[-local-ports {<Number>|<StartingNumber>-<EndingNumber>}]` - Local Ports

This optional parameter specifies logical port associated with the local address to be protected by this policy. The port defaults to any port ('0-0' or '0') but a single port may be specified ('port number' or 'port number-port number').

`[-remote-ports {<Number>|<StartingNumber>-<EndingNumber>}]` - Remote Ports

This optional parameter specifies logical port associated with the remote address to be protected by this policy. The port defaults to any port ('0-0' or '0') but a single port may be specified ('port number' or 'port number-port number').

`[-protocols {<Protocol Number>|<Protocol Name>}]` - Protocols

This optional parameter specifies the protocol to be protected by by this policy. The protocol defaults to any protocol ('any' or '0') but a single protocol may be specified ('tcp', 'udp' or protocol number).

`[-action <IPsec Action Type>]` - Action

This optional parameter specifies the action to be performed when a packet meets the traffic selectors described by this policy. The possible values are 'ESP_TRA' (Apply ESP Transport Mode), 'ESP_UDP' (Apply ESP UDP Encapsulation), 'DISCARD' (discard matching traffic), or 'BYPASS' (send matching traffic in cleartext (not protected by IPsec)). NOTE: If the action is 'BYPASS' or 'DISCARD' and an authentication method is provided, it will be ignored. The default value is 'ESP_TRA'.

[-cipher-suite <Cipher Suite Type>] - Cipher Suite

This optional parameter specifies the suite of algorithms that will be used to protect the traffic. The possible values are:

SUITEB_GCM256: Suite-B-GCM-256 cipher suite as specified in RFC6379.

SUITEB_GMAC256: Suite-B-GMAC-256 cipher suite as specified in RFC6379.

SUITE_AESCBC: Suite consisting of AES256 CBC and SHA512 for ESP and AES256-SHA512-MODP4096 for IKE.

The default value is 'SUITEB_GCM256'.

[-ike-lifetime <integer>] - IKE Security Association Lifetime

This optional parameter specifies the lifetime of an IKE Security Association (in seconds). Shortly before the expiration of the IKE-lifetime, a new IKE security association will be created and the existing IKE security association (and child IPsec security associations) will be destroyed. The default value is 86400 seconds.

[-ipsec-lifetime <integer>] - IPsec Security Association Lifetime

This optional parameter specifies the lifetime of an IPsec Security Association (in seconds). Shortly before the expiration of the ipsec-lifetime, a new IPsec security association will be created and the existing IPsec security association will be destroyed. The default value is 28800 seconds.

[-ipsec-lifetime-bytes <integer>] - IPsec Security Association Lifetime (bytes)

This optional parameter specifies the byte lifetime of an IPsec Security Association. Shortly before the expiration of the ipsec-lifetime-bytes (ipsec-lifetime-bytes have been processed by the IPsec security association), a new IPsec security association will be created and the existing IPsec security association will be destroyed. The default value is 0, i.e infinity bytes.

[-is-enabled {true|false}] - Is Policy Enabled

This optional parameter specifies whether the IPsec policy is enabled or not. Any policy that is created is stored in a replicated database. The 'is-enabled' parameter determines if the policy will be included in those evaluated when determining the best-matched policy to match the traffic selectors of the packet. The default value is 'true'.

[-local-identity <text>] - Local Identity

This optional parameter specifies the local IKE endpoint’s identity for authentication purpose. If this field is not explicitly specified, local-ip-subnet will assume the role for identity. If this field is set to "ANYTHING", then it will be translated to the strongSwan "%any" special identity.

[-remote-identity <text>] - Remote Identity

This optional parameter specifies the remote IKE endpoint’s identity for authentication purpose. If this field is not explicitly specified, remote-ip-subnet will assume the role for identity. If this field is set to "ANYTHING", then it will be translated to the strongSwan "%any" special identity.

[-auth-method <IKE Authentication Method>] - Authentication Method

This optional parameter specifies the authentication method for an IPsec policy. The default value is 'PSK', the pre-shared key authentication method.

[-cert-name <text>] - Certificate for Local Identity

This parameter specifies the certificate name and is mandatory for an IPsec policy using the PKI authentication method. The certificate should have already been installed using security certificate install

command.

Examples

This is an example of the creation of an IPsec policy that protects matching traffic, with all parameters specified. The preshared key can be string of length 18-128 bytes, a sequence hexadecimal digits beginning with 0x or a sequence of Base64 encoded binary data with 0s.

```
cluster-1::> security ipsec policy create -vserver vs_data1 -name Policy1
-local-ip-subnets 192.168.10.1/32 -remote-ip-subnets 192.168.20.1/32
-local-ports 4000 -remote-ports 5001 -protocols tcp -action ESP_TRA
-shared-key This_is_a_shared_key_for_ipsec_policy -ike-version 2 -cipher
-suite SUITEB_GCM256 -ike-lifetime 4000 -ipsec-lifetime 1800 -ipsec
-lifetime-bytes 104880 -is-enabled true

Enter the preshared key for IPsec Policy "Policy1" on Vserver "vs_data1":
Re-enter the preshared key:
```

This is an example of the creation of an IPsec policy that protects matching traffic, with some parameters specified (others will be using the default values). PKI authentication method . is used. In this example, remote-identity does not matter, as long as a trusted certificate is provided.

```
cluster-1::> security ipsec policy create -vserver vs_data1 -name Policy2
-local-ip-subnets 192.168.10.1/32 -remote-ip-subnets 0.0.0.0/0 -local
-ports 2049 -auth-method PKI -cert-name lifcert -local-identity
"CN=lif1_certificate.netapp.com" -remote-identity ANYTHING
```

This is an example of the creation of an IPsec policy that discards matching traffic:

```
cluster-1::> security ipsec policy create -vserver vs_data1 -name
DiscardTraffic -local-ip-subnets 192.168.10.1/32 -remote-ip-subnets
192.168.20.1/32 -action DISCARD
```

Related Links

- security certificate install

security ipsec policy delete

Delete an IPsec policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command deletes an existing IPsec policy.

Parameters

`-vserver <vserver name>` - Vserver

Specifies the Vserver to which the policy belongs. If there is only a single Vserver capable of supporting IPsec, the Vserver parameter is implied.

`-name <text>` - Policy Name

This required parameter specifies the name of the policy to be deleted. The name may be a text string (1-64 bytes), a hexadecimal string (beginning with '0x') or a base64 encoded binary string (beginning with '0s').

Examples

This is an example of IPsec policy deletion where two or more Vservers are capable of supporting IPsec:

```
cluster-1::> security ipsec policy delete -vserver vs_data1 -name
DiscardTraffic
```

This is an example of IPsec policy deletion where only a single Vserver is capable of supporting IPsec:

```
cluster-1::> security ipsec policy delete -name policy1
```

This is an example of an attempt to delete a non-existent IPsec policy:

```
cluster-1::> security ipsec policy delete -vserver vs_data1 -name Discard

Error: There are no entries matching your query.
```

security ipsec policy modify

Modify an IPsec policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command modifies an existing IPsec policy. You cannot modify the name or vserver of a policy. Moving a policy from one Vserver to another or renaming a policy requires that the existing policy be deleted and then a new policy created in the desired Vserver with the desired name.

It is highly recommended that the user set the field `-is-enabled` to `false` prior to making any other modifications to the policy. This will disable the policy and allow all existing IPsec and IKE Security Associations associated with policy to get flushed. Then, the user can modify the policy with the desired changes, along with setting the `-is-enabled` field to `true` to re-enable the policy.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver to which the policy belongs. If there is only a single Vserver capable of supporting IPsec, the Vserver parameter is implied.

-name <text> - Policy Name

This required parameter specifies the name of the policy which may be a text string (1-64 bytes), a hexadecimal string (beginning with '0x') or a base64 encoded binary string (beginning with '0s').

[-local-ip-subnets <IP Address/Mask>,…] - Local IP Subnets

This parameter specifies the IPv4 or IPv6 subnet (address and mask, can be subnet or individual address) representing the local address (range) to be protected by this policy.

[-remote-ip-subnets <IP Address/Mask>,…] - Remote IP Subnets

This parameter specifies the IPv4 or IPv6 subnet (address and mask, can be subnet or individual address) representing the remote address (range) to be protected by this policy.

[-local-ports {<Number>|<StartingNumber>-<EndingNumber>}] - Local Ports

This parameter specifies the logical port associated with the local address to be protected by this policy. The value may be specified by 'port number' or 'port number-port number'.

[-remote-ports {<Number>|<StartingNumber>-<EndingNumber>}] - Remote Ports

This parameter specifies the logical port associated with the remote address to be protected by this policy. The value may be specified by 'port number' or 'port number-port number'.

[-protocols {<Protocol Number>|<Protocol Name>}] - Protocols

This parameter specifies the protocol to be protected by this policy. The protocol may be specified as 'tcp', 'udp' or protocol number.

[-cipher-suite <Cipher Suite Type>] - Cipher Suite

This parameter specifies the suite of algorithms that will be used to protect the traffic. The possible values are:

SUITEB_GCM256: Suite-B-GCM-256 cipher suite as specified in RFC6379.

SUITEB_GMAC256: Suite-B-GMAC-256 cipher suite as specified in RFC6379.

SUITE_AESCBC: Suite consisting of AES256 CBC and SHA512 for ESP and AES256-SHA512-MODP4096 for IKE.

The default value is 'SUITEB_GCM256'.

[-ike-lifetime <integer>] - IKE Security Association Lifetime

This parameter specifies the lifetime of an IKE Security Association (in seconds). Shortly before the expiration of the IKE-lifetime, a new IKE security association will be created and the existing IKE security association (and child IPsec security associations) will be destroyed.

[-ipsec-lifetime <integer>] - IPsec Security Association Lifetime

This parameter specifies the lifetime of an IPsec Security Association (in seconds). Shortly before the expiration of the ipsec-lifetime, a new IPsec security association will be created and the existing IPsec security association will be destroyed.

[-ipsec-lifetime-bytes <integer>] - IPsec Security Association Lifetime (bytes)

This parameter specifies the byte lifetime of an IPsec Security Association. Shortly before the expiration of the ipsec-lifetime-bytes (ipsec-lifetime-bytes have been processed by the IPsec security association), a new IPsec security association will be created and the existing IPsec security association will be destroyed.

[-is-enabled {true|false}] - Is Policy Enabled

This parameter specifies the whether the IPsec policy is enabled or not. Any policy which is created is stored in a replicated database. The 'is-enabled' parameter determines if the policy will be included in those evaluated when determining the best-matched policy to match the traffic selectors of the packet. The default value is 'true'.

[-local-identity <text>] - Local Identity

This optional parameter specifies the local IKE endpoint's identity for authentication purpose. If this field is not explicitly specified, local-ip-subnet will assume the role for identity. If this field is set to "ANYTHING", then it will be translated to the strongSwan "%any" special identity.

[-remote-identity <text>] - Remote Identity

This optional parameter specifies the remote IKE endpoint's identity for authentication purpose. If this field is not explicitly specified, remote-ip-subnet will assume the role for identity. If this field is set to "ANYTHING", then it will be translated to the strongSwan "%any" special identity.

[-cert-name <text>] - Certificate for Local Identity

This optional parameter specifies the certificate name for an IPsec policy using PKI authentication method.

Examples

The following example modifies the local-ip-subnets value of an IPsec policy:

```
cluster-1::> security ipsec policy modify -vserver vs_data1 -name Policy1
-local-ip-subnets 192.168.30.2/32
```

security ipsec policy show

Display IPsec policies

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security ipsec policy show` command displays information about configured IPsec policies. All parameters are optional. This command is supported only when IPsec is enabled.

Running the command with the `-vserver` parameter displays all policies associated with the specified vserver.

You can specify additional parameters to display only information that matches those parameters. For example, to display policies associated with a certain local ip subnet, run the command with the `-local-ip-subnets` parameter.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command displays only the specified fields. Notice that key fields are always displayed.

| [-instance] }

If you specify the `-instance` parameter, the command displays all fields of the policies.

[-vserver <vserver name>] - Vserver

If you specify this parameter, only policies associated with this Vserver will be displayed.

[-name <text>] - Policy Name

This parameter specifies the policy to be displayed.

[-local-ip-subnets <IP Address/Mask>,…] - Local IP Subnets

If you specify this parameter, information about local-ip-subnets will be displayed.

[-remote-ip-subnets <IP Address/Mask>,…] - Remote IP Subnets

If you specify this parameter, information about remote-ip-subnets will be displayed.

[-local-ports {<Number>|<StartingNumber>-<EndingNumber>}] - Local Ports

If you specify this parameter, information about local-ports will be displayed.

[-remote-ports {<Number>|<StartingNumber>-<EndingNumber>}] - Remote Ports

If you specify this parameter, information about remote-ports will be displayed.

[-protocols {<Protocol Number>|<Protocol Name>}] - Protocols

If you specify this parameter, information about protocols will be displayed.

[-action <IPsec Action Type>] - Action

If you specify this parameter, information about action will be displayed.

[-cipher-suite <Cipher Suite Type>] - Cipher Suite

If you specify this parameter, information about cipher-suite will be displayed.

[-ike-lifetime <integer>] - IKE Security Association Lifetime

If you specify this parameter, information about ike-lifetime will be displayed.

[-ipsec-lifetime <integer>] - IPsec Security Association Lifetime

If you specify this parameter, information about ipsec-lifetime will be displayed.

[-ipsec-lifetime-bytes <integer>] - IPsec Security Association Lifetime (bytes)

If you specify this parameter, information about ipsec-lifetime-bytes will be displayed.

[-is-enabled {true|false}] - Is Policy Enabled

If you specify this parameter, information about is-enabled will be displayed.

[-local-identity <text>] - Local Identity

If you specify this parameter, information about local IKE endpoint's identity, if configured, will be displayed.

[-remote-identity <text>] - Remote Identity

If you specify this parameter, information about remote IKE endpoint's identity, if configured, will be displayed.

[-auth-method <IKE Authentication Method>] - Authentication Method

If you specify this parameter, the authentication method of the policy will be displayed.

[-cert-name <text>] - Certificate for Local Identity

If you specify this parameter, the name of the certificate will be displayed.

Examples

The this example displays all policies in all Vservers:

```
cluster-1::> security ipsec policy show
        Policy                                          Cipher
Vserver Name       Local IP Subnet    Remote IP Subnet   Suite
Action
------- ---------- ------------------ ------------------ --------------
-------
vs_data1
        Policy1    192.168.10.1/32    192.168.20.1/32    SUITEB_GCM256
ESP_TRA
        Policy3    192.158.10.10/32   192.158.10.20/32   SUITEB_GCM256
DISCARD
vs_data2
        Policy2    10.10.10.10/32     20.20.20.20/32     SUITE_AESCBC
ESP_TRA
3 entries were displayed.
```

This example displays all of the IPsec policies from a single Vserver:

```
cluster-1::> security ipsec policy show -vserver vs_data1
        Policy                                          Cipher
Vserver Name       Local IP Subnet    Remote IP Subnet   Suite
Action
------- ---------- ------------------ ------------------ --------------
-------
vs_data1
        Policy1    192.168.10.1/32    192.168.20.1/32    SUITEB_GCM256
ESP_TRA
        Policy3    192.158.10.10/32   192.158.10.20/32   SUITEB_GCM256
DISCARD
2 entries were displayed.
```

This example displays a specific policy:

```
cluster-1::> security ipsec policy show -vserver vs_data1 -name Policy1
Vserver Name: vs_data1
                              Policy Name: Policy1
                          Local IP Subnets: 192.168.10.1/32
                         Remote IP Subnets: 192.168.20.1/32
                               Local Ports: 0-0
                              Remote Ports: 0-0
                                 Protocols: any
                                    Action: ESP_TRA
                              Cipher Suite: SUITEB_GCM256
          IKE Security Association Lifetime: 10800
        IPsec Security Association Lifetime: 3600
IPsec Security Association Lifetime (bytes): 0
                         Is Policy Enabled: true
                            Local Identity:
                           Remote Identity:
```

This example displays a specific field from all policies:

```
cluster-1::> security ipsec policy show -fields local-ip-subnets
vserver  name    local-ip-subnets
-------- ------- ----------------
vs_data1 Policy1 192.168.10.1/32
vs_data1 Policy3 192.158.10.10/32
vs_data2
         Policy2 10.10.10.10/32
3 entries were displayed.
```

security key-manager commands

security key-manager delete-key-database

(DEPRECATED)-Deletes the key hierarchy for the Onboard Key Manager

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

> This command is deprecated and might be removed in a future release. Use security key-manager onboard disable instead.

The `security key-manager delete-key-database` command permanently deletes the Onboard Key Manager configuration from all nodes of the cluster.

Examples

The following example deletes the Onboard Key Manager configuration from all nodes of the cluster:

```
cluster-1::*> security key-manager delete-key-database

Warning: This command will permanently delete all keys from the Onboard
Key Manager.
Do you want to continue? {y|n}: y
```

Related Links

- security key-manager onboard disable

security key-manager delete-kmip-config

(DEPRECATED)-Deletes the KMIP configuration

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

> This command is deprecated and may be removed in a future release. Use security key-manager external disable instead.

The `security key-manager delete-kmip-config` command permanently deletes the Key Management Interoperability Protocol (KMIP) server configuration from all nodes of the cluster.

> The keys stored by the external KMIP servers cannot be deleted by ONTAP, and must be deleted by using external tools.

Examples

The following example deletes the KMIP-server configuration from all nodes of the cluster:

```
cluster-1::*> security key-manager delete-kmip-config

Warning: This command will permanently delete the KMIP-server
configuration
         from all nodes of the cluster.
Do you want to continue? {y|n}: y
The KMIP-server configuration has been deleted from all nodes of the
cluster.
The keys stored by the external KMIP servers cannot be deleted by ONTAP,
and must be deleted by using external tools.
```

Related Links

- security key-manager external disable

security key-manager prepare-to-downgrade

Prepares all configured Key managers for downgrade

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command is deprecated and might be removed in a future release.

The `security key-manager prepare-to-downgrade` command disables the Onboard Key Manager features that are not supported in releases prior to ONTAP 9.1.0. The features that are disabled are Onboard Key Manager support for MetroCluster configurations and Volume Encryption (VE).

Examples

The following example disables the Onboard Key Manager support for MetroCluster configurations and Volume Encryption (VE):

```
cluster1::*> security key-manager prepare-to-downgrade
```

security key-manager setup

(DEPRECATED)-Configure key manager connectivity

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

> This command is deprecated and might be removed in a future release. To set up external key manager, use security key-manager external enable , and to set up the Onboard Key Manager use security key-manager onboard enable instead.

The `security key-manager setup` command enables you to configure key management. ONTAP supports two mutually exclusive key management methods: external via one or more key management interoperability protocol (KMIP) servers, or internal via an Onboard Key Manager. This command is used to configure an external or internal key manager. When configuring an external key management server, this command records networking information on all node that is used during the boot process to retrieve keys needed for booting from the KMIP servers. For the Onboard Key Manager, this command prompts you to configure a passphrase to protect internal keys in encrypted form.

This command can also be used to refresh missing onboard keys. For example, if you add a node to a cluster that has the Onboard Key Manager configured, you will run this command to refresh the missing keys.

For the Onboard Key Manager in a MetroCluster configuration, if the security key-manager update-passphrase command is used to update the passphrase on one site, then run the `security key-manager setup` command with the new passphrase on the partner site before proceeding with any key-manager operations.

Parameters

`[-node <nodename>]` - Node Name

This parameter is used only with the Onboard Key Manager when a refresh operation is required (see command description). This parameter is ignored when configuring external key management and during the initial setup of the Onboard Key Manager.

`[-cc-mode-enabled {yes|no}]` - Enable Common Criteria Mode?

When configuring the Onboard Key Manager, this parameter is used to specify that Common Criteria (CC) mode should be enabled. When CC mode is enabled, you will be required to provide a cluster passphrase that is between 64 and 256 ASCII character long, and you will be required to enter that passphrase each time a node reboots.

`[-sync-metrocluster-config {yes|no}]` - Sync MetroCluster Configuration from Peer

When configuring the Onboard Key Manager in a MetroCluster configuration, this parameter is used to indicate that the `security key-manager setup` command has been performed on the peer cluster, and that the `security key-manager setup` command on this cluster should import the peer's configuration.

`[-are-unencrypted-metadata-volumes-allowed-in-cc-mode {yes|no}]` - Are Unencrypted Metadata Volumes Allowed in CC-Mode

If Common Criteria (CC) mode is enabled this parameter allows unencrypted metadata volumes to exist. These metadata volumes are created internally during normal operation. Examples are volumes created during SnapMirror and Vserver migrate operations. The default value is *`no`* .

Examples

The following example creates a configuration for external key management:

```
cluster-1::> security key-manager setup
Welcome to the key manager setup wizard, which will lead you through
the steps to add boot information.

Enter the following commands at any time
"help" or "?" if you want to have a question clarified,
"back" if you want to change your answers to previous questions, and
"exit" if you want to quit the key manager setup wizard. Any changes
you made before typing "exit" will be applied.

Restart the key manager setup wizard with "security key-manager setup". To
accept a default or omit a question, do not enter a value.

Would you like to configure the Onboard Key Manager? {yes, no} [yes]: no
Would you like to configure the KMIP server environment? {yes, no} [yes]:
yes
```

The following example creates a configuration for the Onboard Key Manager:

```
cluster-1::> security key-manager setup
Welcome to the key manager setup wizard, which will lead you through
the steps to add boot information.

Enter the following commands at any time
"help" or "?" if you want to have a question clarified,
"back" if you want to change your answers to previous questions, and
"exit" if you want to quit the key manager setup wizard. Any changes
you made before typing "exit" will be applied.

Restart the key manager setup wizard with "security key-manager setup". To
accept a default or omit a question, do not enter a value.

Would you like to configure the Onboard Key Manager? {yes, no} [yes]: yes
Enter the cluster-wide passphrase for the Onboard Key Manager. To continue
the
configuration, enter the passphrase, otherwise type "exit":
Re-enter the cluster-wide passphrase:
After configuring the Onboard Key Manager, save the encrypted
configuration data
in a safe location so that you can use it if you need to perform a manual
recovery
operation. To view the data, use the "security key-manager backup show"
command.
```

The following example creates a configuration for the Onboard Key Manager with Common Criteria mode enabled:

```
cluster-1::> security key-manager setup -cc-mode-enabled yes
Welcome to the key manager setup wizard, which will lead you through
the steps to add boot information.

Enter the following commands at any time
"help" or "?" if you want to have a question clarified,
"back" if you want to change your answers to previous questions, and
"exit" if you want to quit the key manager setup wizard. Any changes
you made before typing "exit" will be applied.

Restart the key manager setup wizard with "security key-manager setup". To
accept a default or omit a question, do not enter a value.

Would you like to configure the Onboard Key Manager? {yes, no} [yes]: yes
Enter the cluster-wide passphrase for the Onboard Key Manager. To continue
the
configuration, enter the passphrase, otherwise type "exit":
Re-enter the cluster-wide passphrase:
After configuring the Onboard Key Manager, save the encrypted
configuration data
in a safe location so that you can use it if you need to perform a manual
recovery
operation. To view the data, use the "security key-manager backup show"
command.
```

Related Links

- security key-manager external enable
- security key-manager onboard enable
- security key-manager update-passphrase

security key-manager show-key-store

(DEPRECATED)-Displays the configured key manager key stores.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command is deprecated. Use the "security key-manager keystore show" command to display all keystore configurations instead.

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.

[-vserver <vserver name>] - Vserver

If you specify this parameter, then the command will list the key manager configured for the given Vserver.

[-key-store <Key Store>] - Key Store

If you specify this parameter, then the command displays only the vservers that have the given key-store configured.

[-state <Key Store state>] - Key Store State

If you specify this parameter, then the command displays only the vservers that have the given state configured.

[-keystore-type <Key Store Type>] - Key Store Type (Azure/AWS etc)

If you specify this parameter, then the command displays only the vservers that have the given keystore-type configured. This parameter is used to specify a particular type of external key manager. If this parameter is specified and 'key-store' is provided as 'onboard', the "security key-manager show-key-store" command will not return any entries.

[-policy <text>] - Key Manager Policy Name

If you specify this parameter, then the command displays only the vservers that have the given policy.

Examples

The following example shows all configured key managers in the cluster. In the example, the admin vserver has the Onboard Key Manager configured and the data vserver "datavs1" has external key management configured:

```
cluster-1::> security key-manager show-key-store

Vserver                 Key Store Key Store Type
----------------------- --------- --------------
cluster-1               onboard   -
datavs1                 external  AKV
```

security key-manager update-passphrase

(DEPRECATED)-Update cluster-wide passphrase

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

> This command is deprecated and might be removed in a future release. Use security key-manager onboard update-passphrase instead.

The `security key-manager update-passphrase` command provides a way to update the cluster-wide passphrase, created initially by running the security key-manager setup command, that is used for the Onboard Key Manager. This command prompts for the existing passphrase, and if that passphrase is correct then the command prompts for a new passphrase.

When the `security key-manager update-passphrase` command is executed in a MetroCluster configuration, then run the security key-manager setup command with the new passphrase on the partner site before proceeding with any key-manager operations. This allows the updated passphrase to be replicated to the partner site.

Parameters

Examples

The following example updates the cluster-wide passphrase used for the Onboard Key Manager:

```
cluster-1::*> security key-manager update-passphrase

Warning: This command will reconfigure the cluster passphrase for the
Onboard
         Key Manager.
Do you want to continue? {y|n}: y

Enter current passphrase:

Enter new passphrase:

Reenter the new passphrase:
Update passphrase has completed. Save the new encrypted configuration data
in
a safe location so that you can use it if you need to perform a manual
recovery
operation. To view the data, use the "security key-manager backup show"
command.
```

Related Links

- security key-manager onboard update-passphrase
- security key-manager setup

security key-manager config modify

Modify key management configuration options

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command modifies the key management configuration options.

Parameters

`[-cc-mode-enabled {true|false}]` - Enable Common Criteria Mode (privilege: advanced)

This parameter modifies the configuration state of the Onboard Key Manager (OKM) Common Criteria (CC) mode. CC mode enforces some of the policies required by the Common Criteria "Collaborative Protection Profile for Full Drive Encryption-Authorization Acquisition" (FDE-AA cPP) and "Collaborative Protection Profile for Full Drive Encryption-Encryption Engine" documents.

`[-health-monitor-polling-interval <integer>]` - Health Monitor Polling Period (in minutes) (privilege: advanced)

This parameter modifies the the polling interval of the keyserver health monitor at the cluster level.

`[-cloud-kms-retry-count <integer>]` - Cloud KMS connection retry count (privilege: advanced)

This parameter modifies the the cloud keymanager connection retry count at the cluster level.

`[-are-unencrypted-metadata-volumes-allowed-in-cc-mode {true|false}]` - Are Unencrypted Metadata Volumes Allowed in Common Criteria Mode (privilege: advanced)

If Common Criteria (CC) mode is enabled this parameter allows unencrypted metadata volumes to exist. These metadata volumes are created internally during normal operation. Examples are volumes created during SnapMirror and Vserver migrate operations. The default value is *`false`* .

Examples

The following command enables Common Criteria mode in the cluster:

```
cluster-1::*> security key-manager config modify -cc-mode-enabled true
```

The following command modifies the keyserver health monitor polling interval to be 30 minutes:

```
cluster-1::*> security key-manager config modify -health-monitor-polling
-interval 30
```

The following command modifies the cloud keymanager connection retry count to 3:

```
cluster-1::*> security key-manager config modify -cloud-kms-retry-count 3
```

security key-manager config show

Display key management configuration options

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command displays the key management configuration options.

The "cc-mode-enabled" option reflects the current configuration state for Common-Criteria (CC) mode for the Onboard Key Manager. CC mode is an operational mode that enforces some of the policies required by the Common Criteria "Collaborative Protection Profile for Full Drive Encryption-Authorization Acquisition" (FDE-AA cPP) and "Collaborative Protection Profile for Full Drive Encryption-Encryption Engine" documents. The feature can be enabled when the Onboard Key Manager is configured using the security key-manager setup command or after the Onboard Key Manager is configured using the security key-manager config modify command.

Examples

The following example displays the state of all key-manager configuration options:

```
cluster-1::*> security key-manager config show
CC-Mode  health-monitor-polling-interval  cloud-kms-retry-count
Enabled  (in minutes)
-------  ------------                    ---------------------
true     30                              0
```

Related Links

- security key-manager setup
- security key-manager config modify

security key-manager external add-servers

Add external key management servers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command adds the key management servers of the given hosts and ports to the given Vserver's external key manager's list of four possible key management servers. When adding key management servers to the external key manager associated with the admin Vserver, you must run the same command specifying the same set of key servers on the peer cluster. When adding key management servers to a data Vserver, you can run the `security key-manager external add-servers` command on the active cluster only, as the command is replicated to the peer cluster. However, you need to ensure that the key management servers specified are reachable from both clusters. This command is not supported if external key management is not enabled for the Vserver. Use this command to add primary key servers. To modify the list of secondary key servers associated with a primary key server, use the security key-manager external modify-server command.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver on which to add the key management servers.

`-key-servers <Hostname and Port>,…` - External Key Management Servers

Use this parameter to specify the list of additional key management servers that the external key manager uses to store keys.

Examples

The following example adds two key management servers to the list of servers used by the external key manager for Vserver cluster-1. The first key management server's hostname is keyserver1.local and is listening on the default port 5696, and the second key management server's IP is 10.0.0.20 and is listening on port 15696:

```
cluster-1::> security key-manager external add-servers -vserver cluster-1
-key-servers keyserver1.local, 10.0.0.20:15696
```

Related Links

- security key-manager external modify-server

security key-manager external create-config

Create an inactive external key manager

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command enables an ONTAP administrator to create an inactive external key manager configuration for the admin Vserver. This facilitates the switching between the Onboard Key Manager (OKM) and an external key manager (EKM) without migrating the keys to a data SVM

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver on which the external key manager configuration is to be created.

`-key-servers <Hostname and Port>,…` - List of External Key Management Servers

Use this parameter to specify a list of up to four key management servers that the external key manager will use to store keys.

`-client-cert <text>` - Name of the Client Certificate

Use this parameter to specify the unique name of the client certificate that the key management servers will use to verify the identity of the ONTAP client.

`-server-ca-certs <text>,…` - Names of the Server CA Certificates

Use this parameter to specify the unique names of server-ca certificates that ONTAP will use to verify the identity of the key management servers.

Examples

The following example creates an inactive external key manager configuration for Vserver cluster-1 with the

configuration name "default". The command includes three key management servers. The first key server's hostname is ks1.local and is listening on port 5696. The second key server's IP address is 10.0.0.10 and is listening on the default port 5696. The third key server's IPv6 address is fd20:8b1e:b255:814e:32bd:f35c:832c:5a09, and is listening on port 5696.

```
cluster-1::> security key-manager external create-config -vserver cluster-
1 -key-servers
ks1.local:5696,10.0.0.10,[fd20:8b1e:b255:814e:32bd:f35c:832c:5a09]:5696
-client-cert AdminVserverClientCert -server-ca-certs
ServerCaCert1,ServerCaCert2
```

security key-manager external disable

Disable external key management

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command disables the external key manager associated with the given Vserver. If the key manager is in use by ONTAP, you cannot disable it. When disabling the external key manager associated with the admin Vserver, you must run the same command on the peer cluster. When disabling the external key manager for a data Vserver, you can run the `security key-manager external disable` command on the active cluster only, as the command is replicated on the peer cluster. This command is not supported when the Onboard Key Manager is enabled for the given Vserver.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver on which the external key manager is to be disabled.

Examples

The following example removes the external key manager for Vserver cluster-1:

```
cluster-1::*> security key-manager external disable -vserver cluster-1
Warning: This command will permanently delete the external key management
 configuration for Vserver "cluster-1".
 Do you want to continue? {y|n}: y
```

security key-manager external enable

Enable external key management

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command enables the external key manager associated with the given Vserver. This command is not supported when a key manager for the given Vserver is already enabled. When enabling the external key manager associated with the admin Vserver, you must run the same command specifying the same set of key servers on the peer cluster. When enabling the external key manager for a data Vserver, you can run the `security key-manager external enable` command on the active cluster only, as the configuration will be replicated on the peer cluster. However, you must ensure that the key management servers specified in the `security key-manager external enable` command are reachable from both clusters. Only primary key servers can be added using this command.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver on which the external key manager is to be enabled.

`-key-servers <Hostname and Port>,…` - List of External Key Management Servers

Use this parameter to specify the list of up to four key management servers that the external key manager uses to store keys.

`-client-cert <text>` - Name of the Client Certificate

Use this parameter to specify the unique name of the client certificate that the key management servers use to ensure the identity of ONTAP.

`-server-ca-certs <text>,…` - Names of the Server CA Certificates

Use this parameter to specify the unique names of server-ca certificates that ONTAP uses to ensure the identify of the key management servers.

`[-policy <text>]` - Key Manager Policy

Use this parameter to specify a specific key manager security policy to be used by this key manager.

Examples

The following example enables the external key manager for Vserver cluster-1. The command includes three key management servers. The first key server's hostname is ks1.local and is listening on port 15696. The second key server's IP address is 10.0.0.10 and is listening on the default port 5696. The third key server's IPv6 address is fd20:8b1e:b255:814e:32bd:f35c:832c:5a09, and is listening on port 1234.

```
cluster-1::> security key-manager external enable -vserver cluster-1 -key
-servers
ks1.local:15696,10.0.0.10,[fd20:8b1e:b255:814e:32bd:f35c:832c:5a09]:1234
-client-cert AdminVserverClientCert -server-ca-certs
ServerCaCert1,ServerCaCert2
```

security key-manager external modify-server

Modify key server properties

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command modifies configuration information for configured key management servers. When modifying a key management server from the external key manager associated with the admin Vserver, you must run the same command specifying the same set of parameters on the peer cluster. When modifying a key management server from a data Vserver, you can run the `security key-manager external modify-server` command on the active cluster only as the command is replicated on the peer cluster. However, if the password associated with a key management server is modified, then you must run the `security key-manager external modify-server` command specifying the same password on the peer cluster as the password is not replicated between clusters. This command is supported only when external key manager has been enabled for the given Vserver.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver on which to modify the key management server configuration.

`-key-server <Hostname and Port>` - External Key Server

Use this parameter to specify the primary key management server for which the command modifies the configuration.

`[-secondary-key-servers <Remote InetAddress>,…]` - Secondary Key Servers

Use this parameter to specify the secondary key management servers that will be members of the set of clustered key servers. When specifying a secondary key server, a port number cannot be associated with the secondary key server.

`[-timeout <integer>]` - Key Server I/O Timeout (privilege: advanced)

Use this parameter to specify the I/O timeout, in seconds, for the selected key management server.

`[-username <text>]` - Authentication User Name (privilege: advanced)

Use this parameter to specify the username with which ONTAP authenticates with the key management server.

`[-create-remove-timeout <integer>]` - Key Server Timeout for Create and Remove

Use this parameter to specify a shorter I/O timeout, in seconds, to be used for create and delete operations for the selected key management server.

Examples

The following example modifies the I/O timeout to 45 seconds for Vserver cluster-1, key server keyserver1.local:

```
cluster-1::> security key-manager external modify-server -vserver cluster-
1 -key-server keyserver1.local -timeout 45
```

The following example modifies the username and passphrase used to authenticate with key server keyserver1.local:

```
cluster-1::> security key-manager external modify-server -vserver cluster-
1 -key-server keyserver1.local -username ksuser
Enter the password:
Reenter the password:
```

The following example modifies the secondary key management servers secondarykeyserver1.local and secondarykeyserver2.local to be in a cluster configuration with the primary key management server keyserver1.local

```
cluster-1::> security key-manager external modify-server -vserver cluster-
1 -key-server keyserver1.local -secondary-key-servers
secondarykeyserver1.local,secondarykeyserver2.local
```

security key-manager external modify

Modify external key management

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command modifies the external key manager configuration associated with the given Vserver. When modifying the external key manager configuration associated with the admin Vserver, you must run the same command specifying the same parameters on the peer cluster. When modifying the external key manager configuration associated with a data Vserver, you can run the `security key-manager external modify` command on the active cluster only as the configuration modifications are replicated on the peer cluster. This command is not supported when external key management is not enabled for the given Vserver.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver on which the key manager to be modified is located.

`[-client-cert <text>]` - Name of the Client Certificate

Use this parameter to modify the name of the client certificate that the key management servers use to ensure the identity of ONTAP. If the keys of the new certificate do not match the keys of the existing certificate, or if the TLS connectivity with key-management servers fails with the new certificate, the operation fails. Running this command in the diagnostic privilege mode ignores failures and allows the command to complete.

`[-server-ca-certs <text>,…]` - Names of the Server CA Certificates

Use this parameter to modify the names of server-ca certificates that ONTAP uses to ensure the identity of the key management servers. Note that the list provided completely replaces the existing list of certificates. If the TLS connectivity with key-management servers fails with the new list of server-ca certificates, the operation fails. Running this command in the diagnostic privilege mode ignores failures and allows the command to complete.

Examples

The following example updates the client certificate used with the key management servers:

```
cluster-1::> security key-manager external modify -vserver cluster-1
-client-cert NewClientCert
```

security key-manager external remove-servers

Remove external key management servers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command removes the key management servers at the given hosts and ports from the given Vserver's external key manager's list of key management servers. If any of the specified key management servers is the sole storage location for any key that is in use by ONTAP, then you are unable to remove the key server. When removing key management servers from the external key manager associated with the admin Vserver, you must run the same command specifying the same set of key servers on the peer cluster. When removing key management servers from a data Vserver, you can run the `security key-manager external remove-servers` command on the active cluster only as the the command is replicated on the peer cluster. This command is not supported when external key management is not enabled for the given Vserver. Use this command is remove primary key servers. To modify the list of secondary key servers associated with a primary key server, use the security key-manager external modify-server command.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver on which the external key manager is to be removed.

`-key-servers <Hostname and Port>,…` - External Key Management Servers

Use this parameter to specify the list of key management servers that you want to remove from the external key manager.

`[-force {true|false}]` - Bypass OOQ Check?

Set this parameter to true to bypass checks for out of quorum nodes.

Examples

The following example removes the key management server keyserver1.local, listening on the default port of 5696 and the key management server at IP 10.0.0.20, listening on port of 15696.

```
cluster-1::*> security key-manager external remove-servers -vserver
cluster-1
 -key-servers keyserver1.local,10.0.0.20:15696
```

Related Links

- security key-manager external modify-server

security key-manager external restore

Restore the key ID pairs from the key management servers.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command retrieves and restores any current unrestored keys associated with the storage controller from the specified key management servers. When restoring keys from the external key manager associated with the admin Vserver, you must run the same command on the peer cluster. When restoring keys from a data Vserver, you can run the `security key-manager external restore` command on the active cluster only as the command is replicated on the peer cluster. This command is not supported when external key management has not been enabled for the Vserver. This command only restores keys from primary key servers.

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>|local}] - Node

This parameter specifies the name of the node that will load unrestored key IDs into its internal key table. If not specified, all nodes retrieve unrestored keys into their internal key table.

[-vserver <vserver name>] - Vserver Name

This parameter specifies the Vserver for which to list the keys. If not specified, this command restores key for all Vservers.

[-key-server <Hostname and Port>] - Key Server

If this parameter is specified, this command restores keys from the key management server identified by the host and port. If not specified, this command restores keys from all available key management servers.

[-key-id <Hex String>] - Key ID

If you specify this parameter, then the command restores only the key IDs that match the specified value.

[-key-tag <text>] - Key Tag

If you specify this parameter, then the command restores only the key IDs that match the specified key-tag. The key-tag for Volume Encryption Keys (VEKs) is set to the UUID of the encrypted volume. If not specified, all key ID pairs for any key tags are restored.

Examples

The following command restores keys that are currently on a key server but are not stored within the key tables on the cluster. One key is missing for vserver clus- ter-1 on node1, and another key is missing for vserver datavs on node1 and node2:

```
cluster-1::> security key-manager external restore
Node: node1
            Vserver: cluster-1
         Key Server: 10.0.0.1:5696

Key ID
 ---------------------------------------------------------------------------
-------
000000000000000002000000000000100a04fc7303d9abd1e0f00896192fa9c3f0000000000
000000
Node: node1
            Vserver: datavs
         Key Server: tenant.keysever:5696
Key ID
 ---------------------------------------------------------------------------
-------
000000000000000002000000000000400a05a7c294a7abc1e0911897132f49c380000000000
000000
Node: node2
            Vserver: datavs
         Key Server: tenant.keysever:5696

Key ID
 ---------------------------------------------------------------------------
-------
000000000000000002000000000000400a05a7c294a7abc1e0911897132f49c380000000000
000000
```

security key-manager external show-status

Show the set of configured external key management servers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays connectivity information between ONTAP nodes and configured external key management servers.

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>|local}] - Node Name

If you specify this parameter, then the command displays the connectivity information for only the given node.

[-vserver <vserver name>] - Vserver Name

If you specify this parameter, then the command displays the key management servers for only the given Vserver.

[-key-server <Hostname and Port>] - Primary Key Server

If you specify this parameter, then the command displays the connectivity information for only the key management servers with the given primary key server host name or IP address listening on the given port.

[-key-server-status {available|not-responding|unknown}] - Overall Key Server Status

If you specify this parameter, then the command displays the connectivity information for only the key management servers with the given status.

[-status-details <text>] - Overall Key Server Status Details

If you specify this parameter, then the command displays the connectivity information for only the key management servers with the given status details.

[-operational-status {true|false}] - Overall Operational Status

If you specify this parameter, then the command displays the connectivity information for only the key management servers with the given operational status.

[-secondary-key-servers <text>,…] - Secondary Key Servers

If you specify this parameter, then the command displays the connectivity information of only the primary key management servers that have the given secondary key management servers.

[-kmip-version <text>] - KMIP version

If you specify this parameter, then the command displays the connectivity information for only the key management servers with the given KMIP protocol version.

[-port <integer>] - Key Server Port

If you specify this parameter, then the command displays the connectivity information for only the key management servers listening on the given port.

[-key-servers <text>,…] - Individual Key Servers

Use this parameter to show the host name or IP address of the key management server(s).

[-roles <text>,…] - Individual Key Server Roles

Use this parameter to show the role of the key management server(s).

[-key-server-status-list {available|not-responding|unknown}] - Individual Key Server Statuses

Use this parameter to show the status of the key management server(s).

[-status-details-list <text>,…] - Individual Key Server Status Details

Use this parameter to show the status details of the key management server(s).

Examples

The following example lists all configured key management servers for all Vservers:

```
cluster-2::*> security key-manager external show-status
Node: sti219-vsim-sr044e
             Vserver: C1_sti219-vsim-sr044e_1711266079
     Key Server Port: 12014
 KMIP is operational: true

Key Server           Role           Server Status     Reason
-------------------  ------------   ---------------   ------
10.234.2.30          primary        available         -
10.237.80.83         secondary      available         -
10.237.80.223        secondary      available         -
Node: sti219-vsim-sr044f
             Vserver: C1_sti219-vsim-sr044e_1711266079
     Key Server Port: 12014
 KMIP is operational: true

Key Server           Role           Server Status     Reason
-------------------  ------------   ---------------   ------
10.234.2.30          primary        available         -
10.237.80.83         secondary      available         -
10.237.80.223        secondary      available         -
2 entries were displayed.
```

security key-manager external show

Show the set of configured external key management servers.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the external key management servers configured on the cluster for a given Vserver. No entries are displayed when external key management is not enabled for the given Vserver. This command displays the primary external key management servers, along with any associated secondary key servers, configured on the cluster for a given Vserver.

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.

[-vserver <vserver name>] - Vserver Name

If you specify this parameter, then the command displays only the key management servers for the given Vserver.

[-key-server <text>] - Key Server Name with port

If you specify this parameter, then the command displays only the key management servers with the given primary key server host name or IP address listening on the given port.

[-client-cert <text>] - Name of the Client Certificate

If you specify this parameter, then the command displays only the key management servers using a client certificate with the given name.

[-server-ca-certs <text>,…] - Names of the Server CA Certificates

If you specify this parameter, then the command displays only the key management servers using server-ca certificates with the given names.

[-timeout <integer>] - Server I/O Timeout

If you specify this parameter, then the command displays only the key management servers using the given I/O timeout.

[-username <text>] - Authentication User Name

If you specify this parameter, then the command displays only the key management servers using the given authentication username.

[-policy <text>] - Security Policy

If you specify this parameter, then the command displays only the key management servers using the given key manager policy.

[-secondary-key-servers <text>,…] - Secondary Key Servers

If you specify this parameter, then the command displays only the key management servers with the given secondary key servers.

[-create-remove-timeout <integer>] - Key Server Timeout for Create and Remove

If you specify this parameter, then the command displays only the key management servers using the given create-remove I/O timeout.

[-enabled {true|false}] - Is Configuration Enabled?

If you specify this parameter, then the command displays only the key management servers that are enabled.

Examples

The following example lists all configured key management servers for all Vservers:

```
cluster-1::> security key-manager external show
Vserver: datavs
       Client Certificate: datavsClientCert
   Server CA Certificates: datavsServerCaCert1, datavsServerCaCert2
          Security Policy: IBM_Key_Lore
                  Enabled: true

Primary Key Server
 ---------------------------------------------------------------------
keyserver.datavs.com:5696
Vserver: cluster-1
       Client Certificate: AdminClientCert
   Server CA Certificates: AdminServerCaCert
          Security Policy:
                  Enabled: true
Primary Key Server
 ---------------------------------------------------------------------
10.0.0.10:1234
    Secondary Servers: ks1.local, ks2.local
fd20:8b1e:b255:814e:32bd:f35c:832c:5a09:1234
ks1.local:1234
4 entries were displayed.
```

The following example lists all configured key management servers with more detail, including timeouts and usernames:

```
cluster-1::> security key-manager external show -instance
Vserver: datavs
       Client Certificate: datavsClientCert
   Server CA Certificates: datavsServerCaCert1, datavsServerCaCert2
       Primary Key Server: keyserver.datavs.com:5696
                  Timeout: 25
                 Username: datavsuser
          Security Policy: IBM_Key_Lore
                  Enabled: true
    Secondary Key Servers:
Vserver: cluster-1
       Client Certificate: AdminClientCert
   Server CA Certificates: AdminServerCaCert
       Primary Key Server: 10.0.0.10:1234
                  Timeout: 25
                 Username:
          Security Policy:
                  Enabled: true
    Secondary Key Servers: ks1.local, ks2.local
Vserver: cluster-1
       Client Certificate: AdminClientCert
   Server CA Certificates: AdminServerCaCert
       Primary Key Server: fd20:8b1e:b255:814e:32bd:f35c:832c:5a09:1234
                  Timeout: 25
                 Username:
          Security Policy:
                  Enabled: true
    Secondary Key Servers:
Vserver: cluster-1
       Client Certificate: AdminClientCert
   Server CA Certificates: AdminServerCaCert
       Primary Key Server: ks1.local:1234
                  Timeout: 45
                 Username:
          Security Policy:
                  Enabled: true
    Secondary Key Servers:
4 entries were displayed.
```

security key-manager external aws check

Show detailed status of the AWS KMS configurations

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command displays the Amazon Web Service (AWS) Key Management Service (KMS) status.

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>|local}] - Node (privilege: advanced)

If this parameter is specified then the command displays only the AWS KMS status for the given node.

[-vserver <Vserver Name>] - Vserver Name (privilege: advanced)

If this parameter is specified then the command displays only the AWS KMS status for the given Vserver.

[-category <Categories for Cloud KMS status check>] - Component (privilege: advanced)

If this parameter is specified then the command displays only the AWS KMS status for the given category.

```
            Category                    Description
            --------                    -----------
            service_reachability        Cloud KMS Reachability
            ekmip_server                Embedded KMIP Server Reachability
            kms_wrapped_key_status      Status of KMS Wrapped Keys On
 Cluster
```

[-status <Status Check>] - Status (privilege: advanced)

If this parameter is specified then the command displays only the AWS KMS status entries matching the given status.

```
            OK
            FAILED
            UNKNOWN
```

[-detail <text>] - Status Details (privilege: advanced)

This field displays a detailed status message, if available.

Examples

The example below displays the status of all components of all AWS KMS instances configured on node vsim1.

```
cluster-1::> security key-manager external aws check -node vsim1
Vserver: vs1
Node: vsim1

Category: service_reachability
              Status: OK

Category: ekmip_server
              Status: OK

Category: kms_wrapped_key_status
              Status: OK
```

security key-manager external aws disable

Disable AWS KMS

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command disables the Amazon Web Service Key Management Service (AWSKMS) associated with the given Vserver. AWSKMS cannot be disabled if it is in use by ONTAP. This command will fail if AWSKMS has not been enabled for the Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver on which the AWSKMS is to be disabled.

Examples

The following example disables the AWSKMS for Vserver v1.

```
cluster-1::>security key-manager external aws disable -vserver v1
```

security key-manager external aws enable

Enable AWS KMS

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command enables the Amazon Web Service Key Management Service (AWSKMS) associated with the given Vserver. An AWS project and AWSKMS must be deployed on the AWS portal prior to running this command. AWSKMS can only be enabled on a data Vserver that doesn’t already have a key manager configured. AWSKMS cannot be enabled in a MetroCluster environment.

Parameters

`-vserver <Vserver Name>` - Vserver

Use this parameter to specify the Vserver on which the AWSKMS is to be enabled.

`-region <text>` - AWS KMS Region

Use this parameter to specify the region of the deployed AWS project.

`-key-id <text>` - AWS Key Id

Use this parameter to specify the key ID of the deployed AWS project.

`[-access-key-id <text>]` - AWS Access Key ID

Use this parameter to specify the access key ID of the deployed AWS project.

`[-encryption-context <text>]` - Additional Layer of Authentication and Logging

Use this parameter to specify the encryption context to satisfy AWS grant constraint if it is configured. The parameter should be in JSON format.

Examples

The following example enables the AWSKMS for Vserver v1. The parameters in the example command identify an Amazon Web Service (AWS) project application deployed on the AWS. The AWS project application has a region "test_na_region", a key ID "test_KEYID", an access key ID "test_accessKeyID" and an encryption context of "{"team": "NVEsecurity"}".

```
cluster-1::*> security key-manager external aws enable -vserver v1 -region
test_na_region -key-id test_KEYID -access-key-id test_accessKeyID
-encryption-context {"team": "NVEsecurity"}

Enter the Amazon Web Service Key Management Service secret access key:
Press <Enter> when done
```

security key-manager external aws rekey-external

Rekey an external key of the Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command replaces the existing AWS KMS key encryption key (KEK) and results in the key hierarchy being protected by the new user specified AWS KMS KEK. Prior to running this command, the user should have already made the necessary changes on the AWS KMS Portal to use the new KEK. Upon successful completion of this command, the internal keys for the given Vserver will be protected by the new AWS KMS KEK.

Parameters

`-vserver <Vserver Name>` - Vserver

This parameter specifies the Vserver for which ONTAP should rekey the AWS KMS KEK

`-key-id <text>` - AWS Key ID

This parameter specifies the key ID of the new AWS KMS KEK that should be used by ONTAP for the provided Vserver. In the case of automatic AWS KMS KEK rotation, the key ID will be the identifier of the user's already existing AWS KMS Customer Managed Key (CMK). In the case of manual AWS KMS KEK rotation, the key ID will be the identifier of the user's new AWS KMS CMK.

Examples

The following command rekeys the AWS KMS KEK for data Vserver vs1 using a new key-id key3.

```
cluster-1::> security key-manager external aws rekey-external -vserver vs1
-key-id  key3
```

security key-manager external aws rekey-internal

Rekey an internal key of the Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command updates the internal Vserver key hierarchy by rekeying the top-level internal key encryption key (KEK). Upon successful completion of the command, all keys in the Vserver key hierarchy will be protected by the new top-level KEK.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

This parameter specifies the Vserver for which ONTAP should rekey the top-level KEK

Examples

The following command rekeys the top-level KEK for data Vserver vs1.

```
cluster-1::> security key-manager external aws rekey-internal -vserver vs1
```

security key-manager external aws restore

Restore missing keys of AWS KMS

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command retrieves and restores any unrestored keys associated with the given Vserver to each node's internal key tables.

Parameters

`-vserver <Vserver Name>` - Vserver

This parameter specifies the Vserver to which the missing keys will be restored.

Examples

The following command restores missing keys for the data Vserver v1 (which has AWSKMS enabled) to the internal key tables on each node in the cluster.

```
cluster-1::> security key-manager external aws restore -vserver v1
```

security key-manager external aws show

Display AWS KMS configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the Amazon Web Service Key Management Service (AWSKMS) configuration for a given Vserver.

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.

[-vserver <Vserver Name>] - Vserver

If you specify this parameter, then the command displays only the AWSKMS configuration for the given Vserver.

[-region <text>] - AWS KMS Region

If you specify this parameter, then the command displays only the AWSKMS configuration with the given region.

[-key-id <text>] - AWS Key ID

If you specify this parameter, then the command displays only the AWSKMS configuration with the given key-id.

[-access-key-id <text>] - AWS Access Key ID

If you specify this parameter, then the command displays only the AWSKMS configuration with the given access key ID.

[-service <text>] - AWS Service Type

If you specify this parameter, then the command displays only the AWSKMS configurations with the given AWS service type.

[-default-domain <text>] - AWS KMS Default Domain

If you specify this parameter, then the command displays only the AWSKMS configurations with the given AWS KMS default domain.

[-state {available|not-responding|unknown}] - AWS KMS Cluster State

If you specify this parameter, then the command displays only the AWSKMS configurations with the given state. The state can be either available or unknown.

[-unavailable-nodes <text>] - Names of Unavailable Nodes

If you specify this parameter, then the command displays only the AWSKMS configurations with the given unavailable-nodes.

[-polling-period <integer>] - Polling period (in minutes)

If you specify this parameter, then the command displays only the AWSKMS configurations with the given polling period.

[-port <integer>] - AWS KMS Port

If you specify this parameter, then the command displays only the AWSKMS configurations with the given AWS KMS port.

[-verify {true|false}] - Verify the AWS KMS Host

If you specify this parameter, then the command displays only the AWSKMS configurations with the given value of the verify flag.

[-verify-host {true|false}] - Verify the AWS KMS Host's Hostname

If you specify this parameter, then the command displays only the AWSKMS configurations with the given value of the verify-host flag.

[-verify-ip {true|false}] - Verify the AWS KMS Host's IP

If you specify this parameter, then the command displays only the AWSKMS configurations with the given value of the verify-ip flag.

[-host <text>] - AWS KMS Host Name

If you specify this parameter, then the command displays only the AWSKMS configurations with the given AWS KMS host name.

[-encryption-context <text>] - Additional Layer of Authentication and Logging

If you specify this parameter, then the command displays only the AWSKMS configurations with the given value of the AWS encryption-context. The parameter should be in JSON format.

Examples

The following example lists all AWSKMS configurations.

```
cluster-1::>security key-manager external aws show
            Vserver: SAMPLE_VSERVER
             Region: SAMPLE_NA_REGION

Access Key Id                                 State
-----------------------------------------     -------
SAMPLE_ACCESS_KEY_ID                          unknown
SAMPLE_ACCESS_KEY_ID_2                        unknown
Unavailable Nodes:                            node1
```

The following example lists the AWSKMS configurations that have the given encryption context of "{"team": "NVEsecurity"}".

```
cluster-1::>security key-manager external aws show -encryption-context
{"team": "NVEsecurity"}
            Vserver: SAMPLE_VSERVER
             Region: SAMPLE_NA_REGION

Access Key Id                                 State
-----------------------------------------     -------
SAMPLE_ACCESS_KEY_ID                          unknown
Unavailable Nodes:                            node1
```

security key-manager external aws update-credentials

Update AWS secret access key and access key ID

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command allows the user to update the secret access key which is used by the Amazon Web Service Key Management Service (AWSKMS) configured for the given Vserver. The secret access key is initially set by running the security key-manager external aws enable command. This command will fail if AWSKMS has not been enabled for the Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the AWSKMS secret access key will be updated.

`-access-key-id <text>` - Access Key ID (privilege: advanced)

Use this parameter to specify the new access key id of the updated credentials.

`[-skip-verify {true|false}]` - Don't verify user credentials (privilege: advanced)

Set this parameter to true to skip verification of the updated credentials.

Examples

The following example updates the AWSKMS secret access key for Vserver v1.

```
cluster-1::> security key-manager external aws update-credentials -vserver
v1

Enter the new secret access key: Press <Enter> when done
```

Related Links

- security key-manager external aws enable

security key-manager external azure check

Show detailed status of the enabled Azure Key Vault configurations

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command displays the Azure Key Vault (AKV) Key Management Service (KMS) status.

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>|local}]` - Node (privilege: advanced)

If this parameter is specified then the command displays only the AKV status for the given node.

`[-vserver <Vserver Name>]` - Vserver Name (privilege: advanced)

If this parameter is specified then the command displays only the AKV status for the given Vserver.

`[-category <Categories for Cloud KMS status check>]` - Component (privilege: advanced)

If this parameter is specified then the command displays only the AKV status for the given category.

```
            Category                    Description
            --------                    -----------
            service_reachability        Cloud KMS Reachability
            ekmip_server                Embedded KMIP Server Reachability
            kms_wrapped_key_status      Status of KMS Wrapped Keys On
Cluster
```

[-status <Status Check>] - Status (privilege: advanced)

If this parameter is specified then the command displays only the AKV status entries matching the given status.

```
            OK
            FAILED
            UNKNOWN
```

[-detail <text>] - Status Details (privilege: advanced)

This field displays the detailed status message, if available.

Examples

The example below displays the status of all components of all AKV KMS configured on the node.

```
cluster-1::> security key-manager external azure check -node vsim1
Vserver: vs1
Node: vsim1

Category: service_reachability
              Status: OK

Category: ekmip_server
              Status: OK

Category: kms_wrapped_key_status
              Status: OK
```

security key-manager external azure create-config

Create an inactive Azure Key Vault configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command creates an Azure Key Vault (AKV) configuration which can be enabled on a Vserver.

Parameters

-vserver <Vserver Name> - Vserver

The name of the Vserver.

-config-name <text> - Configuration name

The name of the configuration.

`-client-id <text>` - Application (Client) ID of Deployed Azure Application

The ID of the client.

`-tenant-id <text>` - Directory (Tenant) ID of Deployed Azure Application

The ID of the tenant.

`-name {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - Deployed Azure Key Vault DNS Name

The DNS name of the deployed AKV .

`-key-id {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - Key Identifier of AKV Key Encryption Key

The ID of the key.

`[-oauth-host <text>]` - Open Authorization Host Name

The hostname of the OAuth server.

`[-authentication-method <AKV Authentication Method>]` - Authentication Method for Azure Application

Use this parameter to specify the authentication method.

Examples

The example below creates a configuration on a node with the following details: Configuration name: sampleConfig, Client ID: client1, Tenant ID: tenant1, Deployed AKV name: https://samplevault.vault.azure.net, Key ID: https://samplevault.vault.azure.net/keys/key1/keyversion, OAuth Host: https://sampleoauth.net, for Vserver vsTest.

```
cluster-1::> security key-manager external azure create-config -config
-name sampleConfig -client-id client1 -tenant-id tenant1 -name
https://samplevault.vault.azure.net -key-id
https://samplevault.vault.azure.net/keys/key1/keyversion -oauth-host
https://sampleoauth.net -vserver vsTest
```

security key-manager external azure disable

Disable Azure Key Vault

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command disables the Azure Key Vault (AKV) associated with the given Vserver and deletes its configuration. If the AKV is in use by ONTAP, it cannot be disabled. This command is not supported if AKV has not been enabled for the Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver on which the AKV is to be disabled.

Examples

The following example disables the AKV for Vserver v1.

```
cluster-1::>security key-manager external azure disable -vserver v1
```

security key-manager external azure enable

Enable Azure Key Vault

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command enables the Azure Key Vault (AKV) associated with the given Vserver and creates a configuration with the name "default". An Azure application and an AKV must be deployed on the Azure portal prior to running this command. This command is not supported for the admin Vserver, or if a key manager for the given data Vserver is already enabled. This command is also not supported in a MetroCluster environment.

Parameters

`-vserver <Vserver Name>` - Vserver

Use this parameter to specify the Vserver on which the AKV is to be enabled.

`-client-id <text>` - Application (Client) ID of Deployed Azure Application

Use this parameter to specify the client (application) ID of the deployed Azure application.

`-tenant-id <text>` - Directory (Tenant) ID of Deployed Azure Application

Use this parameter to specify the tenant (directory) ID of the deployed Azure application.

`-name {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - Deployed Azure Key Vault DNS Name

Use this parameter to specify the DNS name of the deployed AKV.

`[-authentication-method <AKV Authentication Method>]` - Authentication Method for Azure Application

Use this parameter to specify either client_secret authentication or certificate authentication for the deployed AKV.

`-key-id {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - Key Identifier of AKV Key Encryption Key

Use this parameter to specify the key identifier of the AKV Key Encryption Key (KEK).

`[-oauth-host <text>]` - Open Authorization Host Name

Use this parameter to specify the host name of the Open Authorization server.

Examples

The following example enables the AKV for Vserver v1. An Azure application with client-id "4a0f9c98-c5aa-4275-abe3-2780cf2801c3", tenant-id "8e21f23a-10b9-46fb-9d50-720ef604be98", client secret (not echoed to the screen for security purposes), OAuth host at 10.12.34.1 and an AKV with DNS name "https://akv-keyvault.vault.azure.net" is deployed on the Azure portal. An AKV KEK with DNS name "https://akv-keyvault.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74" is created on the Azure portal for the AKV.

```
cluster-1::>security key-manager external azure enable -client-id
4a0f9c98-c5aa-4275-abe3-2780cf2801c3 -tenant-id 8e21f23a-10b9-46fb-9d50-
720ef604be98 -name https://akv-keyvault.vault.azure.net -key-id
https://akv-
keyvault.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74
-authentication-method client_secret -vserver v1 -oauth-host 10.12.34.1

Enter the client secret for Azure Key Vault:

Re-enter the client secret for Azure Key Vault:
```

The following example enables the AKV for Vserver v1. An Azure application with client-id "4a0f9c98-c5aa-4275-abe3-2780cf2801c3", tenant-id "8e21f23a-10b9-46fb-9d50-720ef604be98", a client certificate (not echoed to the screen for security purposes), OAuth host at 10.12.34.1 and an AKV with DNS name "https://akv-keyvault.vault.azure.net" is deployed on the Azure portal. An AKV KEK with DNS name "https://akv-keyvault.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74" is created on the Azure portal for the AKV.

```
cluster-1::>security key-manager external azure enable -client-id
4a0f9c98-c5aa-4275-abe3-2780cf2801c3 -tenant-id 8e21f23a-10b9-46fb-9d50-
720ef604be98 -name https://akv-keyvault.vault.azure.net -key-id
https://akv-
keyvault.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74
-authentication-method certificate -vserver v1 -oauth-host 10.12.34.1

Enter the client certificate for Azure Key Vault:
```

security key-manager external azure rekey-external

Rekey an external key of the Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command results in the key hierarchy being protected by the user designated AKV key encryption key (KEK). Prior to running this command, the user should have already made the necessary change on the Azure portal to use a new KEK for their key vault. The key-id used in this command is the key ID associated with the user's new AKV KEK. Upon successful completion of this command, the internal keys for the given Vserver will

be protected by the new AKV KEK.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

This parameter specifies the Vserver for which ONTAP should rekey the AKV KEK.

`-key-id {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - Key Identifier of a new AKV Key Encryption Key (privilege: advanced)

This parameter specifies the key id of the new AKV KEK that should be used by ONTAP for the provided Vserver.

Examples

The following command rekeys AKV KEK for data Vserver v1 using a new key, key2 with version 12345678123412341234123456789012.

```
cluster-1::> security key-manager external azure rekey-external -vserver
v1 -key-id  https://kmip-akv-
keyvault.vault.azure.net/keys/key2/12345678123412341234123456789012
```

security key-manager external azure rekey-internal

Rekey an internal key of the Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command rekeys the internal Vserver key hierarchy by changing the top-level internal key encryption key (KEK). Upon successful completion of the command, all keys in the Vserver key hierarchy will be protected by the new top-level KEK.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

This parameter specifies the Vserver for which ONTAP should rekey the SVM KEK.

Examples

The following command rekeys the SVM KEK for data Vserver v1.

```
cluster-1::> security key-manager external azure rekey-internal -vserver
v1
```

security key-manager external azure restore

Restore missing keys of Azure Key Vault

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command retrieves and restores any current unrestored keys associated with the given Vserver to the nodes internal key tables. This command is not supported when AKV has not been enabled for the Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver

This parameter specifies the Vserver to which the missing keys will be restored.

Examples

The following command restores missing keys for the data vserver v1 (which has AKV configuration) to the internal key tables on the cluster.

```
cluster-1::> security key-manager external azure restore -vserver v1
```

security key-manager external azure show

Display Azure Key Vaults configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the Azure Key Vault (AKV) configuration for a given Vserver.

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.

`[-vserver <Vserver Name>]` - Vserver

If you specify this parameter, then the command displays only the AKV configuration for the given Vserver.

`[-config-name <text>]` - Configuration Name

If you specify this parameter, then the command displays only the AKV configurations with the given configuration name.

`[-enabled {true|false}]` - Is This Azure Key Vault Configuration Enabled?

If you specify this parameter, then the command displays only the AKV configurations with the given enabled value.

[-client-id <text>] - Application (Client) ID of Deployed Azure Application

If you specify this parameter, then the command displays only the AKV configuration with the given client id.

[-tenant-id <text>] - Directory (Tenant) ID of Deployed Azure Application

If you specify this parameter, then the command displays only the AKV configuration with the given tenant id.

[-name {scheme://(hostname|IPv4 Address|['IPv6 Address']')…}] - Deployed Azure Key Vault DNS Name

If you specify this parameter, then the command displays only the AKV configuration with the given key vault name.

[-state {available|not-responding|unknown}] - Azure Key Vault Cluster State

If you specify this parameter, then the command displays only the AKV configuration with the given state. The state can be either available or unknown.

[-key-id {scheme://(hostname|IPv4 Address|['IPv6 Address']')…}] - Key Identifier of AKV Key Encryption Key

If you specify this parameter, then the command displays only the AKV configuration with the given key id.

[-unavailable-nodes <text>] - Names of Unavailable Nodes

If you specify this parameter, then the command displays only the AKV configuration with the given unavailable-nodes.

[-authentication-method <AKV Authentication Method>] - AKV Authentication Method

If you specify this parameter, then the command displays only the AKV configurations with the given authentication method.

Examples

The following example lists all Vservers with AKV configuration.

```
cluster-1::>security key-manager external azure show
          Vserver: vs0
      Config Name: default
          Enabled: true
        Client ID: client1
        Tenant ID: tenant1
           Key ID:
https://vault.azure.net/keys/key1/9a2b8c1f7e6d5430e5f4a1b69d87c2e2
   Authentication: client_secret
             Name: https://vault.azure.net
            State: unknown
Unavailable Nodes: node-1,node-2
```

security key-manager external azure update-client-secret

(DEPRECATED)-Update client secret for Azure Key Vault

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

> This command is deprecated and may be removed in a future release. Use security key-manager external azure update-credentials instead.

This command provides a way to update the client secret that is used for the Azure Key Vault (AKV) configured for the given Vserver. The command is initially set by running the security key-manager external azure enable command. This command is not supported if AKV has not been enabled for the Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the AKV client secret is to be updated.

Examples

The following example updates the AKV client secret for the data Vserver v1.

```
cluster-1::> security key-manager external azure update-client-secret
-vserver v1

Enter new client secret:

Re-enter new client secret:
```

Related Links

- security key-manager external azure update-credentials
- security key-manager external azure enable

security key-manager external azure update-credentials

Update client credentials for an Azure Key Vault configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command provides a way to update the authentication credentials that are used for an Azure Key Vault (AKV) configuration. The -config-name parameter can optionally be set to specify a specific AKV configuration of the given Vserver. If the -config-name parameter is not set, the credentials of the given Vserver's enabled configuration are updated. The credentials are initially set by running either the security key-manager external azure enable or security key-manager external azure create-config command.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the AKV client credentials are to be updated.

`[-config-name <text>]` - Configuration Name (privilege: advanced)

Use this parameter to specify the configuration name for which the AKV client credentials are to be updated.

`-authentication-method <AKV Authentication Method>` - Authentication Method for the Azure Application (privilege: advanced)

Use this parameter to specify the authentication method.

Examples

The following examples show different ways of updating the AKV client credentials for the data Vserver v1.

```
cluster-1::> security key-manager external azure update-credentials
-vserver v1 -authentication-method client_secret

Enter new client secret:

Re-enter new client secret:

cluster-1:> security key-manager external azure update-credentials
-vserver v1 -authentication-method certificate

Enter the client certificate for Azure Key Vault:

cluster-1:> security key-manager external azure update-credentials
-vserver v1 -config-name config10 -authentication-method certificate

Enter the client certificate for Azure Key Vault:
```

Related Links

- security key-manager external azure enable
- security key-manager external azure create-config

security key-manager external gcp check

Show detailed status of the Google Cloud KMS configurations

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command displays the Google Cloud Key Management Service (KMS) status.

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>|local}] - Node (privilege: advanced)

If this parameter is specified then the command displays only the Google Cloud KMS status for the given node.

[-vserver <Vserver Name>] - Vserver Name (privilege: advanced)

If this parameter is specified then the command displays only the Google Cloud KMS status for the given Vserver.

[-category <Categories for Cloud KMS status check>] - Component (privilege: advanced)

If this parameter is specified then the command displays only the Google Cloud KMS status for the given category.

```
            Category                    Description
            --------                    -----------
            service_reachability        Cloud KMS Reachability
            ekmip_server                Embedded KMIP Server Reachability
            kms_wrapped_key_status      Status of KMS Wrapped Keys On
 Cluster
```

[-status <Status Check>] - Status (privilege: advanced)

If this parameter is specified then the command displays only the Google Cloud KMS status entries matching the given status.

```
            OK
            FAILED
            UNKNOWN
```

[-detail <text>] - Status Details (privilege: advanced)

This field displays the detailed status message, if available.

Examples

The example below displays the status of all components of all Google Cloud KMS configured on the node.

```
cluster-1::> security key-manager external gcp check -node vsim1
Vserver: vs1
Node: vsim1

Category: service_reachability
              Status: OK

Category: ekmip_server
              Status: OK

Category: kms_wrapped_key_status
              Status: OK
```

security key-manager external gcp disable

Disable a Google Cloud KMS

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command disables the Google Cloud Key Management Service (GCKMS) associated with the given Vserver. GCKMS cannot be disabled if it is in use by ONTAP. This command will fail if GCKMS has not been enabled for the Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver on which the GCKMS is to be disabled.

Examples

The following example disables the GCKMS for Vserver v1.

```
cluster-1::>security key-manager external gcp disable -vserver v1
```

security key-manager external gcp enable

Create and enable a Google Cloud KMS configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command enables the Google Cloud Key Management Service (GCKMS) associated with the given Vserver. A GCP project and GCKMS must be deployed on the GCP portal prior to running this command. GCKMS can only be enabled on a data Vserver that doesn’t already have a key manager configured. GCKMS cannot be enabled in a MetroCluster environment.

Parameters

`-vserver <Vserver Name>` - Vserver

Use this parameter to specify the Vserver on which the GCKMS is to be enabled.

`-project-id <text>` - Google Cloud KMS Project (Application) ID

Use this parameter to specify the project ID of the deployed GCP project.

`-key-ring-name <text>` - Google Cloud KMS Key Ring Name

Use this parameter to specify the key ring name of the deployed GCP project.

`-key-ring-location <text>` - Google Cloud KMS Key Ring Location

Use this parameter to specify the location of the key ring.

`-key-name <text>` - Google Cloud KMS Key Encryption Key Name

Use this parameter to specify the key name of the GCKMS Key Encryption Key (KEK).

`[-port <integer>]` - Google Cloud KMS Port Number

Use this parameter to specify the port of the deployed Google Cloud KMS.

`[-cloudkms-host <text>]` - Google Cloud KMS Host's Subdomain

Use this parameter to specify the Google Cloud KMS Host's Subdomain.

`[-verify {true|false}]` - Verify Identity of Google Cloud KMS?

Use this parameter to specify whether to verify the identity of Google Cloud KMS.

`[-verify-host {true|false}]` - Verify Identity of Google Cloud KMS's Hostname?

Use this parameter to specify whether to verify the identity of Google Cloud KMS hostname.

`[-verify-ip {true|false}]` - Verify Identity of Google Cloud KMS's IP Address?

Use this parameter to specify whether to verify the identity of Google Cloud KMS ip address.

`[-proxy-type {http|https}]` - Proxy Type

Use this parameter to specify the proxy type.

`[-proxy-host <text>]` - Proxy Host

Use this parameter to specify the proxy hostname.

`[-proxy-port <integer>]` - Proxy Port

Use this parameter to specify the proxy port.

`[-proxy-username <text>]` - Proxy Username

Use this parameter to specify the proxy username.

`[-proxy-password <text>]` - Proxy Password

Use this parameter to specify the proxy password.

`[-oauth-host <text>]` - Google Cloud KMS Authorization Host

Use this parameter to specify the host name of the Open Authorization server.

[-oauth-url <text>] - Google Cloud KMS Authorization Url

Use this parameter to specify the URL of the Open Authorization access token.

[-timeout <integer>] - Google Cloud Platform Connection Timeout in Seconds

Use this parameter to specify the Google Cloud connection timeout in seconds.

[-privileged-account <text>] - Google Cloud Privileged Service Account

Use this parameter to specify a privileged service account (email address) with both cloudkms.cryptoKeyVersions.useToEncrypt and cloudkms.cryptoKeyVersions.useToDecrypt permissions. If this parameter is specified, any calls made to the GCKMS will first use iam.serviceAccounts.getAccessToken to impersonate the privileged account.

Examples

The following example enables the GCKMS for Vserver v1. The parameters in the example command identify a Google Cloud Platform (GCP) project application deployed on the GCP. The GCP project application has a Project ID "test_project", a key ring name "key_ring_for_test_project", a key ring location "secure_location_for_key_ring", a key name "testKEK" and OAuth server at 10.12.34.1.

```
cluster-1::*> security key-manager external gcp enable -vserver v1
-project-id test_project -key-ring-name key_ring_for_test_project -key
-ring-location secure_location_for_key_ring -key-name testKEK -oauth-host
10.12.34.1

Enter the contents of the Google Cloud Key Management Service account key
file (json file): Press <Enter> when done
```

security key-manager external gcp rekey-external

Rekey an external key of the Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command replaces the existing GCP key encryption key (KEK) and results in the key hierarchy being protected by the user specified GCP KEK. The GCP key ring in use by the GCP Portal should be updated to use the new KEK prior to running this command. Upon successful completion of this command, the internal keys for the given Vserver will be protected by the new GCP KEK.

Parameters

-vserver <Vserver Name> - Vserver

This parameter specifies the Vserver for which ONTAP should rekey the GCP KEK.

-key-name <text> - Google Cloud KMS Key Encryption Key Name

This parameter specifies the key name of the new GCP KEK that should be used by ONTAP for the provided Vserver.

[-project-id <text>] - Google Cloud KMS Project (Application) ID

This parameter specifies the new project ID of the new GCP KEK that should be used by ONTAP for the provided Vserver.

[-key-ring-name <text>] - Google Cloud KMS Key Ring Name

This parameter specifies the new key ring name of the new GCP KEK that should be used by ONTAP for the provided Vserver.

[-key-ring-location <text>] - Google Cloud KMS Key Ring Location

This parameter specifies the new key ring location of the new GCP KEK that should be used by ONTAP for the provided Vserver.

Examples

The following command rekeys GCP KEK for data Vserver v1 using a new key-name key1.

```
cluster-1::> security key-manager external gcp rekey-external -vserver v1
-key-name key1
```

security key-manager external gcp rekey-internal

Rekey an internal key of the Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command rekeys the internal Vserver key hierarchy by changing the SVM key encryption key (KEK). Upon successful completion of the command, all keys in the Vserver key hierarchy will be protected by the new top-level KEK.

Parameters

-vserver <Vserver Name> - Vserver (privilege: advanced)

This parameter specifies the Vserver for which ONTAP should rekey the SVM KEK.

Examples

The following command rekeys the SVM KEK for data Vserver v1.

```
cluster-1::> security key-manager external gcp rekey-internal -vserver v1
```

security key-manager external gcp restore

Restore missing keys of a Google Cloud KMS

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command retrieves and restores any unrestored keys associated with the given Vserver to each node's internal key tables.

Parameters

`-vserver <Vserver Name>` - Vserver

This parameter specifies the Vserver to which the missing keys will be restored.

Examples

The following command restores missing keys for the data Vserver v1 (which has GCKMS enabled) to the internal key tables on each node in the cluster.

```
cluster-1::> security key-manager external gcp restore -vserver v1
```

security key-manager external gcp show

Display Google Cloud KMS configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the Google Cloud Key Management Service (GCKMS) configuration.

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.

`[-vserver <Vserver Name>]` - Vserver

If you specify this parameter, then the command displays only the GCKMS configurations for the given Vserver.

`[-project-id <text>]` - Google Cloud KMS Project (Application) ID

If you specify this parameter, then the command displays only the GCKMS configurations with the given project id.

`[-key-ring-name <text>]` - Google Cloud KMS Key Ring Name

If you specify this parameter, then the command displays only the GCKMS configurations with the given key ring name.

[-key-ring-location <text>] - Google Cloud KMS Key Ring Location

If you specify this parameter, then the command displays only the GCKMS configurations with the given key ring location.

[-key-name <text>] - Google Cloud KMS Key Encryption Key Name

If you specify this parameter, then the command displays only the GCKMS configurations with the given key name.

[-state {available|not-responding|unknown}] - Google Cloud KMS Cluster State

If you specify this parameter, then the command displays only the GCKMS configurations with the given state. The state can be either available or unknown.

[-unavailable-nodes <text>] - Names of Unavailable Nodes

If you specify this parameter, then the command displays only the GCKMS configurations with the given unavailable-nodes.

[-privileged-account <text>] - Google Cloud Privileged Service Account

If you specify this parameter, then the command displays only the GCKMS configurations with the given privileged account.

[-caller-account <text>] - Google Cloud Caller Service Account

If you specify this parameter, then the command displays only the GCKMS configurations with the given caller account.

Examples

The following example lists all Vservers with GCKMS configuration.

```
cluster-1::>security key-manager external gcp show
            Vserver: SAMPLE_VSERVER
         Project ID: SAMPLE_PROJECT_ID
  Key Ring Location: SAMPLE_KEY_RING_LOCATION
           Key Name: SAMPLE_KEY_NAME
     Caller Account: SAMPLE@CALLER.COM
 Privileged Account: SAMPLE@PRIVILEGED.COM

Key Ring Name                                  State
-----------------------------------------   -------
SAMPLE_KEY_RING_NAME                           unknown
Unavailable Nodes:                             node1
```

security key-manager external gcp update-credentials

Update Google Cloud Project's Service Account Credentials

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command allows the user to update the application credential which is used by the Google Cloud Key Management Service (GCKMS) configured for the given Vserver. The application credential is initially set by running the security key-manager external gcp enable command. This command will fail if GCKMS has not been enabled for the Vserver.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the GCKMS application credential will be updated.

Examples

The following example updates the GCKMS application credential for the data Vserver v1.

```
cluster-1::> security key-manager external gcp update-credentials -vserver
v1

Enter the new application credential: Press <Enter> when done
```

Related Links

- security key-manager external gcp enable

security key-manager health policy modify

Modify the health monitor policy of the given keystore-type

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command modifies the health monitor policy

Parameters

`-keystore-type <Key Store Type>` - Keystore Type

Use this parameter to specify the keystore-type of the health monitor policy to be updated.

`[-enabled {true|false}]` - Health Monitor Enabled?

Use this parameter to specify if the health monitor is enabled.

`[-manage-volume-offline {true|false}]` - Automatically Manage Volume Offline?

Use this parameter to specify if the health monitor can automatically offline the volumes if the key manager cannot be reached.

Examples

The following example updates the AWS 'keystore-type' health monitor policy "enabled" and "manage-volume-

offline" values.

```
cluster-1::> security key-manager health policy modify -keystore-type AWS
-enabled false -manage-volume-offline false
```

security key-manager health policy show

Show the health monitor policy of the given keystore-type

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays the health monitor policies

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.

[-keystore-type <Key Store Type>] - Keystore Type

If you specify this parameter, then the command displays only the health monitor policies with the given keystore-type value.

[-enabled {true|false}] - Health Monitor Enabled?

If you specify this parameter, then the command displays only the health monitor policies with the given enabled value.

[-manage-volume-offline {true|false}] - Automatically Manage Volume Offline?

If you specify this parameter, then the command displays only the health monitor policies with the given manage-volume-offline value.

Examples

The following example lists all health monitor policies.

```
cluster-1::> security key-manager health policy show
              Health Monitor Automatically Manage
Keystore Type Enabled?       Volume Offline?
------------- -------------- --------------------
OKM           true           false
KMIP          true           false
AKV           true           true
GCP           true           false
AWS           true           false
IKP           true           true
6 entries were displayed.
```

security key-manager key create

Create a new authentication key

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command creates a new authentication key (AK) and stores it on the the admin Vserver's key management servers. The command fails if the configured key management servers are already storing more than 256 AKs. If this command fails because there are more than 256 AKs in the cluster, delete unused keys on the Vserver's key management servers and retry the command. This command is not supported when external key management is not enabled for the admin Vserver.

Parameters

`[-key-tag <text>]` - Key Tag

This parameter specifies the key tag to associate with the new authentication key (AK). The default value is the node name. This parameter can be used to help identify created authentication keys (AKs). For example, the security key-manager key query command's key-tag parameter can be used to query for a specific key-tag value.

`[-prompt-for-key {true|false}]` - Prompt for Authentication Passphrase

If you specify this parameter as true, then the command prompts you to enter an authentication passphrase manually instead of generating it automatically. For security reasons, the authentication passphrase you entered is not displayed at the command prompt. You must enter the authentication passphrase a second time for verification. To avoid errors, copy and paste authentication passphrases electronically instead of entering them manually. ONTAP saves the resulting authentication key/key ID pair automatically on the configured key management servers.

Examples

The following example creates an authentication key with the node name as the default key-tag value:

```
cluster-1::> security key-manager key create
Key ID:
0000000000000000020000000000100d0f7c2462d626b739fe81b89f29a092f0000000000
000000
```

The following example creates an authentication key with a user-specified authentication passphrase:

```
cluster-1::> security key-manager key create -prompt-for-key true
Enter a new passphrase:
Reenter the passphrase:
Key ID:
0000000000000000020000000000001006268333f870860128fbe17d393e5083b0000000000
000000
```

Related Links

- security key-manager key query

security key-manager key delete

Delete an existing authentication key

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command removes an authentication key from the configured key management servers on the admin Vserver. The command fails if the given key is currently in use by ONTAP. This command is not supported when external key management is not enabled for the admin Vserver.

Parameters

`-key-id <Hex String>` - Authentication Key ID (privilege: advanced)

Use this parameter to specify the key ID of the key that you want to remove.

Examples

The following example deletes an authentication key:

```
cluster-1::*> security key-manager key delete -key-id
0000000000000000020000000000001006268333f870860128fbe17d393e5083b0000000000
000000
```

security key-manager key migrate

Migrate keys from the admin Vserver's the Onboard Key Manager to a data Vserver's

external key manager and vice versa

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command provides a mechanism to migrate the existing keys of a data Vserver from the admin Vserver's key manager to their own key manager or vice versa. The keys stay the same and the data is not rekeyed, only the keys are migrated from one Vserver's key manager to another. After a successful migration to the new key manager, the data Vserver keys are deleted from the previous key manager.

Parameters

`-from-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the name of the Vserver whose key manager the keys are migrated from.

`-to-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the name of the Vserver whose key manager the keys are migrated to.

Examples

The following example migrates the keys of "datavs" data Vserver from "cluster-1" admin Vserver's key manager to "datavs" data Vserver's key manager:

```
cluster-1::> security key-manager key migrate -from-vserver cluster-1 -to
-vserver datavs
```

The following example migrates the keys of "datavs" data Vserver from "datavs" data Vserver's key manager to "cluster-1" admin Vserver's key manager:

```
cluster-1::> security key-manager key migrate -from-vserver datavs -to
-vserver cluster-1
```

security key-manager key query

Display the key IDs.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the IDs of the keys that are stored in the configured key managers. This command does not update the key tables on the node. Primary key servers, along with any associated secondary key servers, are displayed in the output.

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>|local}] - Node

Use this parameter to specify the name of the node that queries the specified key management servers. If this parameter is not specified, then all nodes query the specified key management servers.

[-vserver <vserver name>] - Vserver Name

Use this parameter to specify the Vserver for which to list the keys.

[-key-server <Hostname and Port>] - Key Server

This parameter specifies the host and port of the key management server that you want to query. This parameter is used only with external key managers.

[-key-id <Hex String>] - Key Identifier

If you specify this parameter, then the command displays only the key IDs that match the specified value.

[-key-tag <text>] - Key Tag

If you specify this parameter, then the command displays only the key IDs that match the specified value. The key-tag for Volume Encryption Keys (VEKs) is set to the UUID of the encrypted volume.

[-key-type <Key Usage Type>] - Key Type

If you specify this parameter, then the command displays only the key IDs that match the specified value.

[-restored {true|false}] - Restored

This parameter specifies whether the key corresponding to the displayed key ID is present in the specified node's internal key table. If you specify 'true' for this parameter, then the command displays the key IDs of only those keys that are present in the system's internal key table. If you specify 'false' for this parameter, then the command displays the key IDs of only those keys that are not present in the system's internal key table.

[-key-store <Key Store>] - Key Store

Use this parameter to specify the key manager type from which to list the keys.

[-key-user <vserver name>] - Key User

If you specify this parameter, then the command displays only the key IDs that are used by the specified Vserver.

[-key-manager <text>] - Key Manager

This parameter specifies the identity of the key manager. For external key managers that will be the host and the port of the key server. In other cases that will be the name of a corresponding key manager.

[-key-store-type <Key Store Type>] - Key Store Type

If you specify this parameter, then the command displays only the key IDs that are used by the specified key manager type.

[-crn <text>] - Cloud Resource Name

This parameter specifies the Cloud Resource Name (CRN) of the key. If you specify this parameter, then the command displays only the key IDs that contains such CRN.

[-policy <text>] - Key Store Policy

This optional parameter specifies the policy name of the key manager. If you specify this parameter, then the command displays only the key IDs that are associated with the specified policy.

[-encryption-algorithm <text>] - Encryption algorithm for the key

This optional parameter specifies the encryption algorithm of the key. If you specify this parameter, then the command displays only the keys of the specified algorithm type.

Examples

The following example shows all of the keys on all configured key servers, and whether or not those keys have been restored for all nodes in the cluster:

```
cluster-1::> security key-manager key query
Node: node1
            Vserver: cluster-1
        Key Manager: onboard
   Key Manager Type: OKM

Key Tag                                  Key Type Encryption   Restored
-------------------------------------  -------- ------------ --------
node1                                    NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000001000c11b3863f78c2273343d7ec5a67762e0000000000
000000
node1                                    NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000001006f4e2513353a674305872a4c9f3bf7970000000000
000000
node1                                    NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000100e1f6b27094485d2d74408bca673b25eb0000000000
000000
node1                                    NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000100ea73be83ec42a7a2bd262f369cda83a40000000000
000000
Node: node1
            Vserver: datavs
        Key Manager: keyserver.datavs.com:5965
   Key Manager Type: KMIP

Key Tag                                  Key Type Encryption   Restored
```

```
------------------------------------  -------- ------------ --------
eb9f8311-e8d8-487e-9663-7642d7788a75  VEK      XTS-AES-256  true
    Key ID:
00000000000000000200000000000004001cb18336f7c8223743d3e75c6a7726e0000000000
000000
9d09cbbf-0da9-4696-87a1-8e083d8261bb  VEK      XTS-AES-256  true
    Key ID:
00000000000000000200000000000004064f2e1533356a470385274a9c3ffb9770000000000
000000
40c3546e-600c-401c-b312-f01be52258dd  VEK      XTS-AES-256  true
    Key ID:
00000000000000000200000000000000401e6f2b09744582d74d084cb6a372be5b0000000000
000000
9b195ecb-35ee-4d11-8f61-15a8de377ad7  VEK      XTS-AES-256  true
    Key ID:
00000000000000000200000000000000040ea73be83ec42a7a2bd262f369cda83a40000000000
000000
Node: node2
            Vserver: cluster-1
        Key Manager: onboard
   Key Manager Type: OKM

Key Tag                                Key Type Encryption   Restored
------------------------------------  -------- ------------ --------
node1                                  NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000001000c11b3863f78c2273343d7ec5a67762e0000000000
000000
node1                                  NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000001006f4e2513353a674305872a4c9f3bf7970000000000
000000
node1                                  NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000000100e1f6b27094485d2d74408bca673b25eb0000000000
000000
node1                                  NSE-AK   AES-256      true
    Key ID:
00000000000000000200000000000000100ea73be83ec42a7a2bd262f369cda83a40000000000
000000
Node: node2
            Vserver: datavs
        Key Manager: keyserver.datavs.com:5965
   Key Manager Type: KMIP

Key Tag                                Key Type Encryption   Restored
```

```
-------------------------------------  -------- ------------ --------
eb9f8311-e8d8-487e-9663-7642d7788a75  VEK      XTS-AES-256  true
    Key ID:
000000000000000002000000000000004001cb18336f7c8223743d3e75c6a7726e0000000000
000000
9d09cbbf-0da9-4696-87a1-8e083d8261bb  VEK      XTS-AES-256  true
    Key ID:
000000000000000002000000000000004064f2e1533356a470385274a9c3ffb9770000000000
000000
40c3546e-600c-401c-b312-f01be52258dd  VEK      XTS-AES-256  true
    Key ID:
000000000000000002000000000000401e6f2b09744582d74d084cb6a372be5b0000000000
000000
9b195ecb-35ee-4d11-8f61-15a8de377ad7  VEK      XTS-AES-256  true
    Key ID:
000000000000000002000000000000040ea73be83ec42a7a2bd262f369cda83a40000000000
000000
```

security key-manager key key-table create

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command creates an entry in the key manager key table. It does not create a new key.

Parameters

`-key-id <Hex String>` - Key ID

This parameter specifies the key ID of the new entry in the table.

`-key-type <Key Usage Type>` - Key Usage Type

This parameter specifies the key type of the new entry. The following key types are supported: NSE-AK, AEK, VEK, NEK and SVM-KEK.

`-encryption-algorithm <text>` - Encryption Algorithm For The Key

This parameter specifies the encryption algorithm associated with the key.

`-creation-time <MM/DD/YYYY HH:MM:SS>` - Key Creation Time

This parameter specifies the date and time that the key was created. The date and time format is "MM/DD/YYYY HH:MM:SS".

Examples

The following example creates an entry in the table:

```
cluster-1::> security key-manager key key-table create -key-id
00000000000000000200000000000500e9ccf3f08e7533d9cd0298e1ebe6c1000000000000
000000 -key-type SVM-KEK -encryption-algorithm AES-256 -creation-time
01/01/2022 01:01:59

cluster-1::> security key-manager key key-table show

Key ID
Key Type Encryption   Creation Time
------------------------------------------------------------------------
------- -------- ------------ --------------------
00000000000000000200000000000500e9ccf3f08e7533d9cd0298e1ebe6c1000000000000
000000  SVM-KEK  AES-256      1/1/2022 01:01:59
1 entry was displayed.
```

security key-manager key key-table delete

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command removes an entry from the key table.

Parameters

`-key-id <Hex String>` - Key ID

Use this parameter to specify the key ID of the entry that you want to remove from the key table.

Examples

The following example deletes an entry from the key table:

```
cluster-1::> security key-manager key key-table show

Key ID
Key Type Encryption   Creation Time
------------------------------------------------------------------------
------- -------- ------------ --------------------
00000000000000000200000000000100239c17902e7515ed397892f75f52e38e0000000000
000000  NSE-AK   AES-256      2/8/2022 10:54:46
00000000000000000200000000000a00a7af571b8397e7df297128fdeb83f4ba0000000000
000000  SVM-KEK  AES-256      1/1/2022 01:01:59
2 entries were displayed.

cluster-1::*> security key-manager key key-table delete -key-id
00000000000000000200000000000100239c17902e7515ed397892f75f52e38e0000000000
000000

cluster-1::> security key-manager key key-table show

Key ID
Key Type Encryption   Creation Time
------------------------------------------------------------------------
------- -------- ------------ --------------------
00000000000000000200000000000a00a7af571b8397e7df297128fdeb83f4ba0000000000
000000  SVM-KEK  AES-256      1/1/2022 01:01:59
1 entry was displayed.
```

security key-manager key key-table modify

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command modifies an entry in the key table. Changes made using this command do not affect the key; only the table entry is modified, not the key itself.

Parameters

`-key-id <Hex String>` - Key ID

This parameter specifies the key ID of the entry to be modified.

`[-key-type <Key Usage Type>]` - Key Usage Type

If this optional parameter is specified, the key type field is modified accordingly.

`[-encryption-algorithm <text>]` - Encryption Algorithm For The Key

If this optional parameter is specified, the encryption algorithm field is modified accordingly.

[-creation-time <MM/DD/YYYY HH:MM:SS>] - Key Creation Time

If this optional parameter is specified, the creation time field is modified accordingly.

Examples

The following example shows the key table before and after the modify command:

```
cluster-1::> security key-manager key key-table show

Key ID
Key Type Encryption   Creation Time
------------------------------------------------------------------------
------- -------- ------------ --------------------
00000000000000000200000000000500e9ccf3f08e7533d9cd0298e1ebe6c1190000000000
000000  VEK      XTS-AES-256  1/1/2022 10:00:00

cluster-1::> security key-manager key key-table modify -key-id
00000000000000000200000000000500e9ccf3f08e7533d9cd0298e1ebe6c1190000000000
000000 -creation-time "12/25/2022 00:00:00"

cluster-1::> security key-manager key key-table show

Key ID
Key Type Encryption   Creation Time
------------------------------------------------------------------------
------- -------- ------------ --------------------
00000000000000000200000000000500e9ccf3f08e7533d9cd0298e1ebe6c1190000000000
000000  VEK      XTS-AES-256  12/25/2022 00:00:00
```

security key-manager key key-table show

Display details of a specific key ID.

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays date and time information for all keys.

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.

[-key-id <Hex String>] - Key ID

If this parameter is specified, the command displays the key that has the specified key ID.

[-key-type <Key Usage Type>] - Key Usage Type

If this parameter is specified, the command only displays information about keys with the specified key type.

[-encryption-algorithm <text>] - Encryption Algorithm For The Key

If this parameter is specified, the command only displays information about keys with the specified encryption algorithm.

[-creation-time <MM/DD/YYYY HH:MM:SS>] - Key Creation Time

If this parameter is specified, the command displays only information about keys with the specified creation time.

Examples

The following example shows all date and time information for all keys:

```
cluster-1::> security key-manager key key-table show

Key ID
Key Type Encryption   Creation Time
--------------------------------------------------------------------------
------- -------- ------------ --------------------
00000000000000000200000000000001000f3ee496cd5820cfb76dd2ce3fa7661b0000000000
000000  NSE-AK   AES-256      1/30/2022 04:21:40
00000000000000000200000000000100658779529aa57ddfef953f305b16c7b20000000000
000000  NSE-AK   AES-256      1/30/2022 04:21:40
00000000000000000200000000000005004100a4355062ea078fdc2fc16b2018d70000000000
000000  VEK      XTS-AES-256  1/30/2022 04:23:14
00000000000000000200000000000a0059f8f7f92612e85664630eed8fb855170000000000
000000  SVM-KEK  AES-256      1/30/2022 04:23:14
4 entries were displayed.
```

security key-manager keystore delete

Remove a disabled key manager keystore configuration.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command is used to delete keystore configurations. Only keystore configurations that are not enabled can be deleted. Use the security key-manager keystore show -enable false command to see a list of all the keystore configurations that are not enabled.

Parameters

`-vserver <Vserver Name>` - Vserver

If you specify this parameter, then the command deletes only the keystore configurations with the given Vserver.

`-keystore-type <Key Store Type>` - Keystore Type

If you specify this parameter, then the command deletes only the keystore configurations with the given keystore type. This command only deletes keystore configurations with a keystore type of Azure Key Vault (AKV).

`-config-name <text>` - Configuration Name

If you specify this parameter, then the command deletes only the keystore configurations with the given configuration name.

Examples

The following example deletes all keystore configurations with the name "NewConfig" that are not enabled.

```
Cluster-1::*> security key-manager keystore delete -vserver * -keystore
-type * -config-name NewConfig
1 entry was deleted.

Cluster-1::*>
```

The following example deletes the keystore configuration on Vserver vs0 with a keystore type AKV and "NewConfig" configuration name.

```
Cluster-1::*> security key-manager keystore delete -vserver vs0 -keystore
-type AKV -config-name NewConfig
1 entry was deleted.

Cluster-1::*>
```

The following example attempts and fails to delete an enabled keystore configuration.

```
Cluster-1::*> security key-manager keystore delete -vserver vs0 -keystore
-type AKV -config-name default

Error: command failed: The keystore configuration with name "default" and
       keystore type "AKV" is currently enabled for Vserver "vs0" and
cannot be
       deleted.

Cluster-1::*>
```

Related Links

- security key-manager keystore show

security key-manager keystore enable

Enable or switch to an inactive keystore configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command enables the keystore represented by the supplied Vserver, keystore-type, and config-name combination. If the specified Vserver already has a keystore enabled, then the enabled configuration is switched to the desired keystore configuration and the previous configuration is left in an inactive/disabled state.

Parameters

`-vserver <vserver name>` - Vserver Name

Use this parameter to specify the Vserver of the desired configuration.

`-keystore-type <Key Store Type>` - Keystore Type

Use this parameter to specify the keystore type of the desired configuration.

`-config-name <text>` - Configuration Name

Use this parameter to specify the configuration name of the desired configuration.

Examples

The following example enables a configuration named sampleConfig for the configured AKV on Vserver vsTest.

```
cluster-1::> security key-manager keystore enable -vserver vsTest
-keystore-type AKV -config-name sampleConfig
```

The following example switches the configuration on Vserver vsTest to a new keystore configuration named newConfig.

```
cluster-1::> security key-manager keystore enable -vserver vsTest
-keystore-type AKV -config-name newConfig
```

security key-manager keystore show

Displays the configured key manager keystores.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays all keystore configurations.

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.

[-vserver <Vserver Name>] - Vserver

If you specify this parameter, then the command displays only the keystore configurations with the given Vserver.

[-keystore-type <Key Store Type>] - Keystore Type

If you specify this parameter, then the command displays only the keystore configurations with the given keystore type.

[-config-name <text>] - Configuration Name

If you specify this parameter, then the command displays only the keystore configurations with the given configuration name.

[-enabled {true|false}] - Configuration Is Enabled?

If you specify this parameter, then the command displays only the keystore configurations with the given enabled value.

[-uuid <UUID>] - Key Store UUID

If you specify this parameter, then the command displays only the keystore configuration with the given universal unique identifier (UUID) value.

[-state <Key Store state>] - Key Store State

If you specify this parameter, then the command displays only the keystores with the given state.

Examples

The following example lists all keystore configurations.

```
Cluster-1::*> security key-manager keystore show

Vserver                 Keystore Type  Config Name        Enabled
----------------------- -------------- ------------------ -------
Cluster-1               OKM            default            true
vs0                     AKV            default            true
vs0                     AKV            AKV-Config-1       false
vs1                     AKV            AWS-Config-1       true
4 entries were displayed.
```

The following example lists only the keystore configurations that are enabled.

```
Cluster-1::*> security key-manager keystore show -enabled true

Vserver                 Keystore Type  Config Name        Enabled
----------------------- -------------- ------------------ -------
Cluster-1               OKM            default            true
vs0                     AKV            default            true
vs1                     AKV            AWS-Config-1       true
3 entries were displayed.
```

security key-manager onboard create-config

Add an inactive Onboard Key Manager

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command enables an ONTAP administrator to create an inactive configuration for the named Vserver. This facilitates the switching between the Onboard Key Manager (OKM) and an external key manager (EKM) without migrating the keys to a data SVM.

Parameters

Examples

The following example creates an inactive configuration for the Onboard Key Manager:

```
cluster-1::*> security key-manager onboard create-config
```

security key-manager onboard disable

Disable the Onboard Key Manager

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command disables the Onboard Key Manager associated with the admin Vserver and permanently deletes the Onboard Key Manager configuration associated with the admin Vserver. The Onboard Key Manager cannot be disabled if there are any encrypted volumes that use encryption keys created by the Onboard Key Manager. This command fails if the Onboard Key Manager is not enabled.

Examples

The following example disables the Onboard Key Manager for the admin Vserver:

```
cluster-1::*> security key-manager onboard disable

Warning: This command will permanently delete all keys from Onboard Key
         Manager.
Do you want to continue? {y|n}: y
```

security key-manager onboard enable

Enable the Onboard Key Manager

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command enables the Onboard Key Manager for the admin Vserver.

Parameters

`[-cc-mode-enabled {yes|no}]` - Enable Common Criteria Mode?

Use this parameter to specify whether the Common Criteria (CC) mode should be enabled or not. When CC mode is enabled, you are required to provide a cluster passphrase that is between 64 and 256 ASCII characters long, and you are required to enter that passphrase each time a node reboots. CC mode cannot be enabled in a MetroCluster configuration.

`[-are-unencrypted-metadata-volumes-allowed-in-cc-mode {yes|no}]` - Are Unencrypted Metadata Volumes Allowed in Common Criteria Mode

If Common Criteria (CC) mode is enabled this parameter allows unencrypted metadata volumes to exist. These metadata volumes are created internally during normal operation. Examples are volumes created during SnapMirror and Vserver migrate operations. the default value is *`no`* .

Examples

The following example enables the Onboard Key Manager for the admin Vserver cluster-1:

```
cluster-1::> security key-manager onboard enable

Enter the cluster-wide passphrase for the Onboard Key Manager:

Re-enter the cluster-wide passphrase:

After configuring the Onboard Key Manager, save the encrypted
configuration data in a safe location so that you can use it if you need
to perform a manual recovery operation. To view the data, use the
"security key-manager onboard show-backup" command.
```

The following example enables the Onboard Key Manager for the admin Vserver cluster-1 with the cluster-wide passphrase "myPassphrase" on an external USB mass storage device:

```
cluster-1::> security key-manager onboard enable -vserver cluster-1 -okm
-on-usb-enabled true -onboard-passphrase myPassphrase

Enter the cluster-wide passphrase for the Onboard Key Manager:

Re-enter the cluster-wide passphrase:
After configuring the Onboard Key Manager, save the encrypted
configuration data in a safe location so that you can use it if you need
to perform a manual recovery operation. To view the data, use the
"security key-manager onboard show-backup" command.
```

security key-manager onboard show-backup

Display the Onboard Key Management backup

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the backup information for the Onboard Key Manager for the admin Vserver, which can be used to recover the cluster in case of catastrophic situations. The information displayed is for the cluster as a whole (not individual nodes).

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.

Examples

The following example displays the Onboard Key Manager backup data for the admin Vserver:

```
cluster-1::> security key-manager onboard show-backup
--------------------------BEGIN BACKUP--------------------------
TmV0QXBwIEtleSBCbG9iAAEAAAAEAAAAcAEAAAAAAADuD+byAAAAACEAAAAAAAAA
QAAAAAAAAABvOlH0AAAAAMh7qDLRyH1DBz12piVdy9ATSFMT0C0TlYFss4PDjTaV
dzRYkLd1PhQLxAWJwOIyqSr8qY1SEBgm1IWgE5DLRqkiAAAAAAAAACgAAAAAAAA
3WTh7gAAAAAAAAAAAAAAAAIAAAAAAAgAZJEIWvdeHr5RCAvHGclo+wAAAAAAAAA
IgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3tJA/
LRzUQRHwv+1aWvAAAAAAAAAACQAAAAAAAAgAAAAAAAACdhTcvAAAAAJ1PXeBf
ml4NBsSyV1B4jc4A7cvWEFY6lLG6hc6tbKLAHZuvfQ4rIbYAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABOZXRBcHAgS2V5IEJs
b2IAAQAAAAMAAAAYAQAAAAAADA5/
ccAAAAAIgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3t
JA/LRzUQRHwv+1aWvAAAAAAAAAACIAAAAAAAAKAAAAAAAAACI8z/
bAAAAAAAAAAAAAAAAAgAAAAAAAQAbxMcI4qiaMS4Uts5tTUnUAAAAAAAAAAAkAAA
AAAAAAIAAAAAAAAAAqwxTcwAAAACkiwBAI3YeeV3jMFg5SmyjLSgoK/
qc8FAmMMcrRXY6uriulnL0WPB/
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBL
ZXkgQmxvYgABAAAAAwAAABgBAAAAAAAA1cNLLwAAAAAiAAAAAAAAACgAAAAAAAA
Q5NxHQAAAAAAAAAAAAAAAIAAAAAAAkAave0kD8tHNRBEfC/
7Vpa8AAAAAAAAAAAIgAAAAAAAAAoAAAAAAAAAJ4/
cQsAAAAAAAAAAAAAAACAAAAAAABAF6JCZch+IF+ZeOutovhv8oAAAAAAAAAACQA
AAAAAAAAgAAAAAAAAAAN3Zq7AAAAALO7qD20+H8TuGgSauEHoqAyWcLv4uA0m2rr
H4nPQM0nrDRYRa9SCv8AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAA
---------------------------END BACKUP---------------------------
```

security key-manager onboard sync

Sync the Onboard Key Manager keys

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command synchronizes missing onboard keys on any node in the cluster. For example, if you add a node to a cluster that has the Onboard Key Manager configured, you should then run this command to synchronize the keys. In a MetroCluster configuration, if the security key-manager onboard enable command is used to enable the Onboard Key Manager on one site, then run the `security key-manager onboard sync` command on the partner site. In a MetroCluster configuration, if the security key-manager onboard update-passphrase command is used to update the passphrase on one site, then run this command with the new passphrase on the partner site before proceeding with any key management operations.

Parameters

Examples

The following example synchronizes the Onboard Key Manager key database across all nodes in the cluster. In a MetroCluster configuration, this command synchronizes nodes in the local site.

```
cluster-1::> security key-manager onboard sync
```

Related Links

- security key-manager onboard enable
- security key-manager onboard update-passphrase

security key-manager onboard update-passphrase

Update the Onboard Key Manager Passphrase

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command provides a way to update the cluster-wide passphrase that is used for the Onboard Key Manager and initially created by running the security key-manager onboard enable command. This command prompts for the existing passphrase, and if that passphrase is correct then the command prompts for a new passphrase. When the Onboard Key Manager is enabled for the admin Vserver, run the security key-manager onboard show-backup command after updating the passphrase and save the output for emergency recovery scenarios. When the `security key-manager onboard update-passphrase` command is executed in a MetroCluster configuration, then run the security key-manager onboard sync command with the new passphrase on the partner site before proceeding with any key-manager operations. This allows the updated passphrase to be replicated to the partner site.

Parameters

Examples

The following example updates the cluster-wide passphrase used for the Onboard Key Manager:

```
cluster-1::*> security key-manager onboard update-passphrase

Warning: This command will reconfigure the cluster passphrase for onboard
         key management.
Do you want to continue? {y|n}: y

Enter current passphrase:

Enter new passphrase:

Reenter the new passphrase:
Update passphrase has completed. Save the new encrypted configuration data
in
a safe location so that you can use it if you need to perform a manual
recovery
operation. To view the data, use the "security key-manager onboard show-
backup"
command.
```

Related Links

- security key-manager onboard enable
- security key-manager onboard show-backup
- security key-manager onboard sync

security key-manager onboard verify-backup

Verify the onboard key management backup and its passphrase

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command verifies the backup data and the passphrase of the Onboard Key Manager for the admin Vserver.

Examples

The following example displays the verification of the onboard key management backup data for the admin Vserver:

```
cluster-1::> security key-manager onboard verify-backup

Enter the Onboard Key Manager backup data:
TmV0QXBwIEtleSBCbG9iAAEAAAAEAAAAcAEAAAAAAADuD+byAAAAACEAAAAAAAAA
QAAAAAAAAABvOlH0AAAAAMh7qDLRyH1DBz12piVdy9ATSFMT0C0TlYFss4PDjTaV
dzRYkLd1PhQLxAWJwOIyqSr8qY1SEBgm1IWgE5DLRqkiAAAAAAAAACgAAAAAAAA
3WTh7gAAAAAAAAAAAAAAAAIAAAAAAAgAZJEIWvdeHr5RCAvHGclo+wAAAAAAAAA
IgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3tJA/
LRzUQRHwv+1aWvAAAAAAAAAACQAAAAAAAAgAAAAAAAACdhTcvAAAAAJ1PXeBf
ml4NBsSyV1B4jc4A7cvWEFY6lLG6hc6tbKLAHZuvfQ4rIbYAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABOZXRBcHAgS2V5IEJs
b2IAAQAAAAMAAAAYAQAAAAAADA5/
ccAAAAAIgAAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAAACAAAAAAAJAGr3t
JA/LRzUQRHwv+1aWvAAAAAAAAAACIAAAAAAAAAKAAAAAAAAACI8z/
bAAAAAAAAAAAAAAAAAgAAAAAAAQAbxMcI4qiaMS4Uts5tTUnUAAAAAAAAAAAkAAA
AAAAAAIAAAAAAAAAAqwxTcwAAAACkiwBAI3YeeV3jMFg5SmyjLSgoK/
qc8FAmMMcrRXY6uriulnL0WPB/
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBL
ZXkgQmxvYgABAAAAAwAAABgBAAAAAAA1cNLLwAAAAAiAAAAAAAAACgAAAAAAAAA
Q5NxHQAAAAAAAAAAAAAAAIAAAAAAAkAave0kD8tHNRBEfC/
7Vpa8AAAAAAAAAAIgAAAAAAAAoAAAAAAAAAJ4/
cQsAAAAAAAAAAAAAAACAAAAAAABAF6JCZch+IF+ZeOutovhv8oAAAAAAAAAACQA
AAAAAAAAgAAAAAAAAAAN3Zq7AAAAALO7qD20+H8TuGgSauEHoqAyWcLv4uA0m2rr
H4nPQM0nrDRYRa9SCv8AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAA

Enter the cluster-wide passphrase for the Onboard Key Manager:

The Onboard Key Manager backup data was successfully verified
```

security key-manager policy show

Display key manager policy

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays the defined key management key policies.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You can use '-fields ?' to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-policy <text>] - Policy name

If you specify this parameter, then the command will list the key manager policy with the given name.

[-check-key-on-online {true|false}] - Pull key from key manager during volume online?

If you specify this parameter, then the command displays only the key manager policies with the given check-key-on-online value.

[-purge-key-on-offline {true|false}] - Purge key from memory during volume offline?

If you specify this parameter, then the command displays only the key manager policies with the given purge-key-on-offline value.

[-support-on-admin-vserver {true|false}] - Support policy on admin Vserver?

If you specify this parameter, then the command displays only the key manager policies with the given support-on-admin-vserver value.

[-key-manager-attribute-required {true|false}] - Key manager attribute required for volume?

If you specify this parameter, then the command displays only the key manager policies with the given key-manager-attribute-required value.

Examples

The following example lists all configured key management policies:

```
cluster-1::*> security key-manager policy show

Policy             Check Key on Online?  Purge Key on Offline?
-----------------  --------------------  ---------------------
IBM_Key_Lore       true                  true
```

security login commands

security login create

Add a login method

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login create` command creates a login method for the management utility. A login method consists of a user name, an application (access method), and an authentication method. A user name can be associated with multiple applications. It can optionally include an access-control role name. If an Active Directory, LDAP, or NIS group name is used, then the login method gives access to users belonging to the specified group. If the user is a member of multiple groups provisioned in the security login table, then the user will get access to a combined list of the commands authorized for the individual groups.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver name of the login method.

`-user-or-group-name <text>` - User Name or Group Name

This specifies the user name or Active Directory, LDAP, or NIS group name of the login method. The Active Directory, LDAP, or NIS group name can be specified only with the *`domain`* or *`nsswitch`* authentication method and *`ontapi`* and *`ssh`* application. If the user is a member of multiple groups provisioned in the security login table, then the user will get access to a combined list of the commands authorized for the individual groups.

`-application <text>` - Application

This specifies the application of the login method. Possible values include amqp, console, http, ontapi, rsh, snmp, service-processor, ssh, and telnet.

Setting this parameter to *`service-processor`* grants the user access to the Service Processor (SP). Because the SP supports only password as the first factor of authentication, when you set this parameter to *`service-processor`* , you must also set the -authentication-method parameter to *`password`* . Vserver user accounts cannot access the SP. Therefore, you cannot use the -vserver parameter when you set this parameter to *`service-processor`* .

`-authentication-method <text>` - Authentication Method

This specifies the authentication method for login. Possible values include the following:

- cert - SSL certificate authentication
- community - SNMP community strings
- domain - Active Directory authentication
- nsswitch - LDAP or NIS authentication
- password - Password
- publickey - Public-key authentication
- usm - SNMP user security model
- saml - SAML authentication

`[-remote-switch-ipaddress <IP Address>]` - Remote Switch IP Address

This specifies the IP address of the remote switch. The remote switch could be a cluster switch monitored by cluster switch health monitor (CSHM) or a Fibre Channel (FC) switch monitored by MetroCluster health monitor (MCC-HM). This parameter is applicable only when the application is *`snmp`* and authentication

method is *usm* (SNMP user security model).

-role <text> - Role Name

This specifies an access-control role name for the login method.

[-comment <text>] - Comment Text

This specifies comment text for the user account, for example, "Guest account". The maximum length is 128 characters.

[-is-ns-switch-group {yes|no}] - Whether Ns-switch Group

This specifies whether *user-or-group-name* is an LDAP or NIS group. Possible values are yes or no. Default value is no.

[-second-authentication-method {none|publickey|password|nsswitch|domain|totp}] - Second Authentication Method

This specifies the authentication method for the login. It will be used as the second factor for authentication. This parameter can be specified for *ssh* , *http* and *service-processor* applications only. For *ssh* application, possible values include the following:

- password - Password
- publickey - Public-key authentication
- nsswitch - NIS or LDAP authentication
- domain - Active Directory authentication
- totp - TOTP authentication
- none - default value

For the *http* and *service-processor* applications, possible values include the following:

- publickey - Public-key authentication
- none - default value

[-is-ldap-fastbind {yes|no}] - LDAP Fastbind Authentication

This flag specifies whether the authentication is LDAP fastbind or Not. Default:false

Examples

The following example illustrates how to create a login that has the user name *monitor* , the application *ssh* , the authentication method *password* , and the access-control role *guest* for Vserver *vs1.netapp.com* :

```
cluster1::> security login create -vserver vs1.netapp.com -user-or-group
-name monitor
  -application ssh -authentication-method password -role guest
```

The following example illustrates how to create a login that has the user name *monitor* , the application *ontapi* , the authentication method *password* , and the access-control role *vsadmin* for Vserver *vs1.netapp.com* :

```
cluster1::> security login create -vserver vs1.netapp.com -user-or-group
-name monitor
  -application ontapi -authentication-method password -role vsadmin
```

The following example illustrates how to create a login that has the user name *monitor* , the application *ssh* , the authentication method *publickey* , and the access-control role *guest* for Vserver *vs1.netapp.com* :

```
cluster1::> security login create -vserver vs1.netapp.com -user-or-group
-name monitor
  -application ssh -authentication-method publickey -role guest
```

The following example illustrates how to create a login that has the user name *monitor* , the application *http* , the authentication method *cert* , and the access-control role *admin* for Vserver *vs1.netapp.com* :

```
cluster1::> security login create -vserver vs1.netapp.com -user-or-group
-name monitor
  -application http -authentication-method cert -role admin
```

The following example illustrates how to create a login that has the Active Directory group name *adgroup* in *DOMAIN1* , the application *ssh* , the authentication method *domain* , and the access-control role *vsadmin* for Vserver *vs1.netapp.com* :

```
cluster1::> security login create -vserver vs1.netapp.com
  -user-or-group-name DOMAIN1\adgroup -application ssh
  -authentication-method domain -role vsadmin
```

The following example illustrates how to create a login that has a group name *nssgroup* in the LDAP or NIS server, the application *ontapi* , the authentication method *nsswitch* , and the access-control role *vsadmin* for Vserver *vs1.netapp.com* . Here *is-ns-switch-group* must be set to *yes* :

```
cluster1::> security login create -vserver vs1.netapp.com -user-or-group
-name nssgroup
  -application ontapi -authentication-method nsswitch -role vsadmin
  -is-ns-switch-group yes
```

The following example illustrates how to create a login that has the user name *monitor* , the application *ssh* , the authentication method *password* , the second authentication method *publickey* and the access-control role *vsadmin* for Vserver *vs1.netapp.com* :

```
cluster1::> security login create -vserver vs1.netapp.com -user-or-group
-name monitor
  -application ssh -authentication-method password
  -second-authentication-method publickey -role vsadmin
```

The following example illustrates how to create a login that has the user name *monitor* , the application *ssh* , the authentication method *password* , the second authentication method *none* and the access-control role *vsadmin* for Vserver *vs1.netapp.com* :

```
cluster1::> security login create -vserver vs1.netapp.com -user-or-group
-name monitor
  -application ssh -authentication-method password
  -second-authentication-method none -role vsadmin
```

The following example illustrates how to create a login that has the user name *spuser* , the application *service-processor* , the authentication method *password* , the second authentication method *publickey* and the access-control role *admin* for the Administrative Vserver *cluster1* :

```
cluster1::> security login create -vserver cluster1 -user-or-group-name
spuser
  -application service-processor -authentication-method password
  -second-authentication-method publickey -role admin
```

security login delete

Delete a login method

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login delete` command deletes a login method.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name of the login method.

`-user-or-group-name <text>` - User Name or Group Name

This specifies the user name or Active Directory, LDAP, or NIS group name of the login method that is to be deleted. A user name can be associated with multiple applications.

`-application <text>` - Application

This specifies the application of the login method. Possible values include amqp, console, http, ontapi, rsh, snmp, service-processor, ssh, and telnet.

`-authentication-method <text>` - Authentication Method

This specifies the authentication method of the login method. Possible values include the following:

- cert - SSL certificate authentication
- community - SNMP community strings
- domain - Active Directory authentication
- nsswitch - LDAP or NIS authentication
- password - Password
- publickey - Public-key authentication
- usm - SNMP user security model
- saml - SAML authentication

`[-remote-switch-ipaddress <IP Address>]` - Remote Switch IP Address

This specifies the IP address of the remote switch. The remote switch could be a cluster switch monitored by cluster switch health monitor (CSHM) or a Fibre Channel (FC) switch monitored by MetroCluster health monitor (MCC-HM). This parameter is applicable only when the application is *`snmp`* and authentication method is *`usm`* (SNMP user security model).

Examples

The following example illustrates how to delete a login that has the username *`guest`* , the application *`ssh`* , and the authentication method *`password`* for Vserver *`vs1.netapp.com`* :

```
cluster1::> security login delete -user-or-group-name guest
  -application ssh -authentication-method password -vserver vs1.netapp.com
```

The following example illustrates how to delete a login that has the username *`guest`* , the application *`ontapi`* , and the authentication method *`cert`* for Vserver *`vs1.netapp.com`* :

```
cluster1::> security login delete -user-or-group-name guest
  -application ontapi -authentication-method cert -vserver vs1.netapp.com
```

The following example illustrates how to delete a login that has the Active Directory group name *`adgroup`* in *`DOMAIN1`* , the application *`ssh`* , and the authentication method *`domain`* for Vserver *`vs1.netapp.com`* :

```
cluster1::> security login delete -user-or-group-name DOMAIN1\adgroup
  -application ssh -authentication-method domain -vserver vs1.netapp.com
```

The following example illustrates how to delete a login that has a group name *`nssgroup`* in the LDAP or NIS server, the application *`ontapi`* , and the authentication method *`nsswitch`* for Vserver *`vs1.netapp.com`* :

```
cluster1::> security login delete -user-or-group-name nssgroup
  -application ontapi -authentication-method nsswitch -vserver
vs1.netapp.com
```

The following example illustrates how to delete a login that has the username *spuser* , the application *service-processor* , the authentication method *password* and the second authentication method *publickey* for the Administrative Vserver *cluster1* :

```
cluster1::> security login delete -user-or-group-name spuser
  -application service-processor -authentication-method password -vserver
cluster1
```

security login expire-password

Expire user's password

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login expire-password` command expires a specified user account password, forcing the user to change the password upon next login.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver to which the user account belongs.

`-username <text>` - Username

This specifies the user name of the account whose password you want to expire.

`[-hash-function {sha512|sha256}]` - Password Hash Function

This optionally specifies the password-hashing algorithm used for encrypting the passwords that you want to expire. The supported values include are as follows:

- *sha512* - Secure hash algorithm (512 bits)
- *sha256* - Secure hash algorithm (256 bits)
- *md5* - Message digest algorithm (128 bits)

`[-lock-after <integer>]` - Lock User Account After N days (privilege: advanced)

This optionally specifies the number of days after which the new password hash policy will be enforced. The enforcement will lock all user accounts that are still compliant with the provided hash algorithm using `-hash-function` parameter.

Examples

The following command expires the password of the 'jdoe' user account which belongs to the 'vs1.netapp.com' Vserver.

```
cluster1::> security login expire-password -vserver vs1.netapp.com
-username jdoe
```

The following command expires all user account passwords that are encrypted with the MD5 hash function.

```
cluster1::> security login expire-password -vserver * -username * -hash
-function md5
```

The following command expires the password of any Vserver's user account named 'jdoe' that is encrypted with the MD5 hash function.

```
cluster1::> security login expire-password -vserver * -username jdoe -hash
-function md5
```

The following command expires the password of the 'vs1.netapp.com' Vserver user account named 'jdoe' that is encrypted with the MD5 hash function.

```
cluster1::> security login expire-password -vserver vs1.netapp.com
-username jdoe -hash-function md5
```

The following command expires all user account passwords that are encrypted with the MD5 hash function and enforce the new password hash policy after 180 days.

```
cluster1::> security login expire-password -vserver * -username * -hash
-function md5 -lock-after 180
```

security login lock

Lock a user account with password authentication method

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login lock` command locks a specified account, preventing it from accessing the management interface. This command only applies to user accounts configured with the password authentication method where the password is set.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver to which the user account belongs.

`-username <text>` - Username

This specifies the user name of the account that is to be locked.

Examples

The following example locks a user account named 'jdoe' which belongs to the Vserver 'vs1.netapp.com'.

```
cluster1::> security login lock -vserver vs1.netapp.com -username jdoe
```

security login modify

Modify a login method

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login modify` command modifies the access-control role name of a login method. If the user is a member of multiple groups provisioned in the security login table, then the user will get access to a combined list of the commands authorized for the individual groups.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver name of the login method.

`-user-or-group-name <text>` - User Name or Group Name

This specifies the user name, Active Directory, LDAP, or NIS group name of the login method that is to be modified. A user name can be associated with multiple applications. If the user is a member of multiple groups provisioned in the security login table, then the user will get access to a combined list of the commands authorized for the individual groups.

`-application <text>` - Application

This specifies the application of the login method. Possible values include amqp, console, http, ontapi, rsh, snmp, service-processor, ssh, and telnet.

`-authentication-method <text>` - Authentication Method

This specifies the authentication method of the login method. Possible values include the following:

- cert - SSL certificate authentication
- community - SNMP community strings
- domain - Active Directory authentication
- nsswitch - LDAP or NIS authentication

- password - Password
- publickey - Public-key authentication
- usm - SNMP user security model
- saml - SAML authentication

[-remote-switch-ipaddress <IP Address>] - Remote Switch IP Address

This specifies the IP address of the remote switch. The remote switch could be a cluster switch monitored by cluster switch health monitor (CSHM) or a Fibre Channel (FC) switch monitored by MetroCluster health monitor (MCC-HM). This parameter is applicable only when the application is *snmp* and authentication method is *usm* (SNMP user security model).

[-role <text>] - Role Name

This modifies the access-control role name for the login method.

[-comment <text>] - Comment Text

This specifies comment text for the user account, for example, "Guest account". The maximum length is 128 characters.

[-is-ns-switch-group {yes|no}] - Whether Ns-switch Group

This specifies if *user-or-group-name* is an LDAP or NIS group. Possible values are yes or no. Default value is no.

[-second-authentication-method {none|publickey|password|nsswitch|domain|totp}] - Second Authentication Method

This specifies the authentication method for the login method. It will be used as the second factor for authentication. This parameter can be specified for *ssh* , *http* and *service-processor* applications. For *ssh* application, possible values include the following:

- password - Password
- publickey - Public-key authentication
- nsswitch - NIS or LDAP authentication
- domain - Active Directory authentication
- totp - TOTP authentication
- none - default value

For the *http* and *service-processor* applications, possible values include the following:

- publickey - Public-key authentication
- none - default value

[-is-ldap-fastbind {yes|no}] - LDAP Fastbind Authentication

This flag specifies whether modify is allowed or not when the authentication is LDAP fastbind.

Examples

The following example illustrates how to modify a login method that has the user name *guest* , the application *ontapi* , and the authentication method *password* to use the access-control role *guest* for Vserver *vs1.netapp.com* :

```
cluster1::> security login modify -user-or-group-name guest
  -application ontapi -authentication-method password -role guest
  -vserver vs1.netapp.com
```

The following example illustrates how to modify a login method that has the user name *guest* , the application *ssh* , and the authentication method *publickey* to use the access-control role *vsadmin* for Vserver *vs1.netapp.com* :

```
cluster1::> security login modify -user-or-group-name guest
  -application ssh -authentication-method publickey -role vsadmin
  -vserver vs1.netapp.com
```

The following example illustrates how to modify a login method that has the group name *nssgroup* , the application *ontapi* , and the authentication method *nsswitch* to use the access-control role *readonly* for Vserver *vs1.netapp.com* . Here *is-ns-switch-group* must be set to *yes* :

```
cluster1::> security login modify -user-or-group-name nssgroup
  -application ontapi -authentication-method nsswitch -role readonly
  -vserver vs1.netapp.com -is-ns-switch-group yes
```

The following example illustrates how to modify a login method that has the user name *guest* , the application *ssh* , and the authentication method *publickey* to use the second-authentication-method *password* for Vserver *vs1.netapp.com* :

```
cluster1::> security login modify -user-or-group-name guest
  -application ssh -authentication-method publickey
  -second-authentication-method password -vserver vs1.netapp.com
```

The following example illustrates how to modify a login method to have individual authentication methods that have the user name *guest* , the application *ssh* , and the authentication method *publickey* to use the second-authentication-method *none* for Vserver *vs1.netapp.com* :

```
cluster1::> security login modify -user-or-group-name guest
  -application ssh -authentication-method publickey
  -second-authentication-method none -vserver vs1.netapp.com
```

The following example illustrates how to modify a login method that has the user name *spuser* , the application *service-processor* , and the authentication method *password* to use the second-authentication-method *publickey* for the Administrative Vserver *cluster1* :

```
cluster1::> security login modify -user-or-group-name spuser
  -application service-processor -authentication-method password
  -second-authentication-method publickey -vserver cluster1
```

security login password-prepare-to-downgrade

Reset password features introduced in the Data ONTAP version

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

If the password of the system administrator is not encrypted with an encryption type supported by releases earlier than ONTAP 9.0, this command prompts the administrator for a new password and encrypt it using a supported encryption type on each cluster or at each site in a MetroCluster configuration. In a MetroCluster configuration, this command must be run on both sites. The password for all other users are marked as "expired". This causes them to be re-encrypted using a compatible encryption type. The expired passwords are changed with an internally generated password. The administrator must change the passwords for all users before the users can login. The users are prompted to change their password upon login. This command disables the logging of unsuccessful login attempts. The command must be run by a user with the cluster admin role from a clustershell session on the console device. This user must be unlocked. If you fail to run this command, the revert process fails.

Parameters

`-disable-feature-set <downgrade version>` - Data ONTAP Version (privilege: advanced)

This parameter specifies the ONTAP version that introduced the password feature set.

Examples

The following command disables the logging of unsuccessful login attempts.

```
cluster1::*> security login password prepare-to-downgrade -disable-feature
-set 8.3.1

         Warning: This command will disable the MOTD feature that prints
unsuccessful login attempts.
         Do you want to continue? {y|n}: y
cluster1::*>
```

The following command prompts system administrator to enter password and encrypt it with the hashing algorithm supported by releases earlier than ONTAP 9.0.

```
cluster1::*> security login password prepare-to-downgrade -disable-feature
-set 9.0.0
Warning: If your password is not encrypted with an encryption type
supported by
                 releases earlier than ONTAP 9.0.0, this command will
prompt you
                 for a new password and encrypt it using a supported
encryption type on
     each cluster or at each site in a MetroCluster configuration. In a
     MetroCluster configuration, this command must be run on both sites.
     The password for all other users are marked as "expired" and
     changed to an internally generated password. The administrator must
change
     the passwords for all users before the users can login. The users are
                 prompted to change their password upon login.
        Do you want to continue? {y|n}:

        Enter a new password:
        Enter it again:

        cluster1::*>
```

security login password

Modify a password for a user

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login password` command resets the password for a specified user. The command prompts you for the user's old and new password.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name of the login method.

`-username <text>` - Username

This optionally specifies the user name whose password is to be changed. If you do not specify a user, the command defaults to the user name you are currently using.

Examples

The following command initiates a password change for the 'admin' user account of the 'vs1.netapp.com' Vserver.

```
cluster1::> security login password -username admin -vserver
vs1.netapp.com
```

security login show

Show user login methods

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login show` command displays the following information about user login methods:

- User name
- Application (amqp, console, http, ontapi, rsh, snmp, service-processor, ssh, or telnet)
- Authentication method (community, password, publickey, or usm)
- Role name
- Whether the account is locked
- Whether the user name refers to *`nsswitch`* group
- Password hash function
- LDAP fastbind authentication

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.

[-vserver <vserver name>] - Vserver

Displays the login methods that match the specified Vserver name.

[-user-or-group-name <text>] - User Name or Group Name

Displays the login methods that match this parameter value. Value can be a user name or Active Directory, LDAP, or NIS group name.

[-application <text>] - Application

Displays the login methods that match the specified application type. Possible values include amqp, console, http, ontapi, rsh, snmp, service-processor, ssh, and telnet.

[-authentication-method <text>] - Authentication Method

Displays the login methods that match the specified authentication method. Possible values include the following:

- cert - SSL certificate authentication
- community - SNMP community strings
- domain - Active Directory authentication
- nsswitch - LDAP or NIS authentication
- password - Password
- publickey - Public-key authentication
- usm - SNMP user security model
- saml - SAML authentication

`[-remote-switch-ipaddress <IP Address>]` - Remote Switch IP Address

Displays the login methods that match the specified IP address of the remote switch. The remote switch could be a cluster switch monitored by cluster switch health monitor (CSHM) or a Fibre Channel (FC) switch monitored by MetroCluster health monitor (MCC-HM). This parameter is applicable only when the application is *`snmp`* and authentication method is *`usm`* (SNMP user security model).

`[-role <text>]` - Role Name

Displays the login methods that match the specified role.

`[-is-account-locked {yes|no}]` - Account Locked

Displays the login methods that match the specified account lock status.

`[-comment <text>]` - Comment Text

Displays the login methods that match the specified comment text.

`[-is-ns-switch-group {yes|no}]` - Whether Ns-switch Group

This specifies whether *`user-or-group-name`* is an LDAP or NIS group. Possible values are yes or no.

`[-hash-function {sha512|sha256}]` - Password Hash Function (privilege: advanced)

Displays the login methods that match the specified password-hashing algorithm. Possible values are:

- *`sha512`* - Secure hash algorithm (512 bits)
- *`sha256`* - Secure hash algorithm (256 bits)
- *`md5`* - Message digest algorithm (128 bits)

`[-second-authentication-method {none|publickey|password|nsswitch|domain|totp}]` - Second Authentication Method

Displays the login methods that match the specified authentication method to be used as the second factor. Presently, *`ssh`* , *`http`* and *`service-processor`* are the only applications supporting a second factor of authentication. For *`ssh`* application, possible values include the following:

- password - Password
- publickey - Public-key authentication
- nsswitch - NIS or LDAP authentication
- domain - Active Directory authentication
- totp - TOTP authentication

- none - default value

For the *`http`* and *`service-processor`* applications, possible values include the following:

- publickey - Public-key authentication
- none - default value

.

`[-is-ldap-fastbind {yes|no}]` - LDAP Fastbind Authentication

Displays the authentication methods that are LDAP fastbind.

Examples

The example below illustrates how to display information about all user login methods:

```
cluster1::> security login show

Vserver: cluster1
                                                                      Second
User/Group                 Authentication                 Acct
Authentication
Name           Application Method        Role Name        Locked Method
-------------- ----------- ------------- ---------------- ------
--------------
admin          amqp        password      admin            no     none
admin          console     password      admin            no     none
admin          http        password      admin            no     none
admin          ontapi      password      admin            no     none
admin          service-processor
                           password      admin            no     none
admin          ssh         password      admin            no     none
autosupport    console     password      autosupport      no     none
user1          ssh         publickey     admin            -      none
user2          ssh         password      admin            no     publickey
spuser         service-processor
                           password      admin            no     publickey

Vserver: vs1.netapp.com
                                                                      Second
User/Group                 Authentication                 Acct
Authentication
Name           Application Method        Role Name        Locked Method
-------------- ----------- ------------- ---------------- ------
--------------
vsadmin        http        password      vsadmin          yes    none
vsadmin        ontapi      password      vsadmin          yes    none
vsadmin        ssh         password      vsadmin          yes    none
12 entries were displayed.
```

security login unlock

Unlock a user account with password authentication method

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login unlock` command unlocks a specified account, enabling it to access the management interface. This command only applies to user accounts configured with the password authentication method where the password is set.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver to which the user account belongs.

`-username <text>` - Username

This specifies the user name of the account that is to be unlocked.

Examples

The following command unlocks a user account named jdoe which belongs to the Vserver vs1.netapp.com.

```
cluster1::> security login unlock -vserver vs1.netapp.com -username jdoe
```

security login whoami

Show the current user, trust score of the user and role of this session

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login whoami` command displays the name, role and the trust score of the user logged in at the current console session. It takes no options or other parameters.

Examples

The following example shows that the current session is logged in by using the 'admin' user account:

```
cluster1::> whoami
               (security login whoami)
User: admin
            Role: admin
     Trust Score: 90
```

security login banner modify

Modify the login banner message

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login banner modify` command modifies the login banner. The login banner is printed just before the authentication step during the SSH and console device login process.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Use this parameter to specify the Vserver whose banner will be modified. Use the name of the cluster admin Vserver to modify the cluster-level message. The cluster-level message is used as the default for data Vservers that do not have a message defined.

{ `[-message <text>]` - Login Banner Message

This optional parameter can be used to specify a login banner message. If the cluster has a login banner message set, the cluster login banner will be used by all data Vservers as well. Setting a data Vserver's login banner will override the display of the cluster login banner. To reset a data Vserver's login banner to use the cluster login banner, use this parameter with the value `"-"`.

If you use this parameter, the login banner cannot contain newlines (also known as end of lines (EOLs) or line breaks). To enter a login banner message with newlines, do not specify any parameter. You will be prompted to enter the message interactively. Messages entered interactively can contain newlines.

Non-ASCII characters must be provided as Unicode UTF-8.

| `[-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}]` - Download URI for the Banner Message }

Use this parameter to specify the URI from where the login banner will be downloaded. Note that the message must not exceed 2048 bytes in length. Non-ASCII characters must be provided as Unicode UTF-8.

Examples

This example shows how to enter a login banner interactively:

```
cluster1::> security login banner modify
Enter the login banner for Vserver "cluster1".
Max size: 2048. Enter a blank line to terminate input. Press Ctrl-C to
abort.
0        1         2         3         4         5         6         7
8
12345678901234567890123456789012345678901234567890123456789012345678901234
567890
Authorized users only!
cluster1::>
```

security login banner show

Display the login banner message

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login banner show` command displays the login banner.

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.

[-vserver <Vserver Name>] - Vserver Name

Selects login banners that match the specified value. Use the name of the admin Vserver to specify the cluster-level login banner.

[-message <text>] - Login Banner Message

Selects login banners that match the specified value. By default, this command will not display unconfigured, or empty, login banners. To display all banners, specify `-message` `*` .

Examples

The following shows sample output from this command:

```
cluster1::> security login banner show
Message
 ---------------------------------------------------------------------------
----
Authorized users only!
cluster1::>
```

security login domain-tunnel create

Add authentication tunnel Vserver for administrative Vserver

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command establishes a gateway (tunnel) for authenticating Windows Active Directory (AD) domain users' access to the cluster.

Before using this command to establish the tunnel, the following must take place:

- You must use the security login create command to create one or more AD domain user accounts that will be granted access to the cluster.
- The -authmethod parameter of the security login create command must be set to 'domain'.
- The -username parameter of the security login create command must be set to a valid AD domain user account that is defined in a Windows Domain Controller's Active Directory. The user account must be specified in the format of `<domainname>\<username>` , where "domainname" is the name of the CIFS domain server.

- You must identify or create a CIFS-enabled data Vserver that will be used for Windows authentication with the Active Directory server. This Vserver is the tunnel Vserver, and it must be running for this command to succeed.

Only one Vserver can be used as the tunnel. If you attempt to specify more than one Vserver for the tunnel, ONTAP returns an error. If the tunnel Vserver is stopped or deleted, AD domain users' authentication requests to the cluster will fail.

Parameters

`-vserver <vserver>` - Authentication Tunnel Vserver

This parameter specifies a data Vserver that has been configured with CIFS. This Vserver will be used as the tunnel for authenticating AD domain users' access to the cluster.

Examples

The following commands create an Active Directory domain user account ('DOMAIN1\Administrator') for the 'cluster1' cluster, create a data Vserver ('vs'), create a CIFS server ('vscifs') for the Vserver, and specify 'vs' as the tunnel for authenticating the domain user access to the cluster.

```
cluster1::> security login create -vserver cluster1 -username
DOMAIN1\Administrator -application ssh -authmethod domain -role admin
    cluster1::> vserver create -vserver vs -rootvolume vol -aggregate aggr
-rootvolume-security-style mixed
    cluster1::> vserver cifs create -vserver vs -cifs-server vscifs
-domain companyname.example.com -ou CN=Computers
    cluster1::> security login domain-tunnel create -vserver vs
```

Related Links

- security login create

security login domain-tunnel delete

Delete authentication tunnel Vserver for administrative Vserver

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login domain-tunnel delete` command deletes the tunnel established by the security login domain-tunnel create command. An error message will be generated if no tunnel exists.

Examples

The following command deletes the tunnel established by security login domain-tunnel create .

```
cluster1::> security login domain-tunnel delete
```

Related Links

- security login domain-tunnel create

security login domain-tunnel modify

Modify authentication tunnel Vserver for administrative Vserver

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login domain-tunnel modify` command modifies or replaces the tunnel Vserver. If a tunnel Vserver is not already specified, it sets the current tunnel Vserver with this Vserver, otherwise, it replaces the current tunnel Vserver with the Vserver that you specify. If the tunnel Vserver is changed, authentication requests via previous Vserver will fail. See security login domain-tunnel create for more information.

Parameters

`[-vserver <vserver>]` - Authentication Tunnel Vserver

This parameter specifies a Vserver that has been configured with CIFS and is associated with a Windows Domain Controller's Active Directory authentication. This Vserver will be used as an authentication tunnel for login accounts so that they can be used with administrative Vservers.

Examples

The following command modifies the tunnel Vserver for administrative Vserver.

```
cluster1::> security login domain-tunnel modify -vserver vs
```

Related Links

- security login domain-tunnel create

security login domain-tunnel show

Show authentication tunnel Vserver for administrative Vserver

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login domain-tunnel show` command shows the tunnel Vserver that was specified by the security login domain-tunnel create or security login domain-tunnel modify command.

Examples

The example below shows the tunnel Vserver, *`vs`* , that is currently used as an authentication tunnel. The output informs you that the table is currently empty if tunnel Vserver has not been specified.

```
cluster1::> security login domain-tunnel show
  Tunnel Vserver: vs
```

Related Links

- security login domain-tunnel create
- security login domain-tunnel modify

security login duo create

Add a Duo Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo create` creates the Duo configuration on the Vserver.

Parameters

`-vserver <vserver name>` - Vserver

This parameter specifies the Vserver name of the Duo configurations.

`[-is-enabled {true|false}]` - Duo Enable Status

This parameter enable the Duo authentication.

`-integration-key <text>` - Duo Intgration Key

This parameter specifies Duo integration key.

`-secret-key <text>` - Duo Secret Key

This parameter specifies Duo secret key.

`-apihost <text>` - Duo API Host

This parameter specifies the Duo API host.

`[-autopush {true|false}]` - Duo Auto Push Config

This optionally specifies the autopush configurations. Configuring this to true will allows user to authenticate with single tap on their mobile device.

`[-pushinfo {true|false}]` - Information with Push

This parameter specifies the push info configurations. Configuring this to true will provide additional information in the push notification, such as the name of the application. This helps users to verify that they are logging in to the correct service and provides an additional layer of security

`[-fail-mode {safe|secure}]` - Duo Fail Safe Config

This parameter specifies the fail mode configurations. Possible values includes safe and secure.

[-http-proxy <text>] - HTTP Proxy URL with Port Number

This parameter specifies the http proxy configurations.

[-max-prompts <integer>] - Duo Max Attempts

This parameter specifies the number of retry before Duo authentication return failure.

[-comment <text>] - Comment

This optionally specifies comment text for the Duo configuration. Note that comment text should be enclosed in quotation marks.

Examples

The following command creates a Duo configurations for Vserver vs1.

```
cluster1::> security login duo create -vserver vs1 -integration-key
AA11A1AAAA1AAAA11A1A -secret-key
xxxxxxxxxx11111111117bd5a3b060947b617355ecf353627c50b1xxxxxxxxxx -apihost
api-99X9X9XX.duosecurity.com
-comment "This is a Duo Config"
```

security login duo delete

Delete a Duo Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo delete` command deletes the Duo configuration.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver name of the Duo configurations.

Examples

The following command deletes the Duo configurations of the Vserver vs1.

```
cluster1::> security login duo delete -vserver vs1
```

security login duo modify

Modify a Duo Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo modify` modifies the Duo configuration.

Parameters

`-vserver <vserver name>` - Vserver

This parameter specifies the Vserver name of the Duo configurations.

`[-is-enabled {true|false}]` - Duo Enable Status

This parameter enable the Duo authentication.

`[-integration-key <text>]` - Duo Intgration Key

This parameter specifies Duo integration key.

`[-secret-key <text>]` - Duo Secret Key

This parameter specifies Duo secret key.

`[-apihost <text>]` - Duo API Host

This parameter specifies the Duo API host.

`[-autopush {true|false}]` - Duo Auto Push Config

This optionally specifies the autopush configurations. Configuring this to true will allows user to authenticate with single tap on their mobile device.

`[-pushinfo {true|false}]` - Information with Push

This parameter specifies the push info configurations. Configuring this to true will provide additional information in the push notification, such as the name of the application. This helps users to verify that they are logging in to the correct service and provides an additional layer of security

`[-fail-mode {safe|secure}]` - Duo Fail Safe Config

This parameter specifies the fail mode configurations. Possible values includes safe and secure.

`[-http-proxy <text>]` - HTTP Proxy URL with Port Number

This parameter specifies the http proxy configurations.

`[-max-prompts <integer>]` - Duo Max Attempts

This parameter specifies the number of retry before Duo authentication return failure.

`[-comment <text>]` - Comment

This parameter specifies comment text for the Duo. Note that comment text should be enclosed in quotation marks.

Examples

The following command modify the Duo configurations for the Vserver vs1.

```
cluster1::> security login duo modify -vserver vs1 -is-enabled false
```

security login duo show

Display Duo Configurations

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo show` command displays information about the Duo configurations.

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.

[-vserver <vserver name>] - Vserver

Selects the Duo configuration Vserver that match this parameter value.

[-is-enabled {true|false}] - Duo Enable Status

Selects the Duo Configurations that match this parameter value.

[-integration-key <text>] - Duo Intgration Key

Selects the Duo Configurations that match this parameter value.

[-fingerprint <text>] - Duo Secret Key Fingerprint

Selects the Duo Configurations that match this parameter value.

[-apihost <text>] - Duo API Host

Selects the Duo Configurations that match this parameter value.

[-autopush {true|false}] - Duo Auto Push Config

Selects the Duo Configurations that match this parameter value.

[-pushinfo {true|false}] - Information with Push

Selects the Duo Configurations that match this parameter value.

[-fail-mode {safe|secure}] - Duo Fail Safe Config

Selects the Duo Configurations that match this parameter value.

[-http-proxy <text>] - HTTP Proxy URL with Port Number

Selects the Duo Configurations that match this parameter value.

[-max-prompts <integer>] - Duo Max Attempts

Selects the Duo Configurations that match this parameter value.

[-status <text>] - Duo Status

Selects the Duo Configurations that match this parameter value.

[-comment <text>] - Comment

Selects the Duo Configurations that match this parameter value.

Examples

The example below displays the Duo information for the Vserver VS1.

```
cluster1::> security login duo show
                    Vserver: VS1
                    Enabled: true
                     Status: OK
            Integration Key: AA11A1AAAA1AAAA11A1A
            SHA Fingerprint:

xxxxxxxxxx11111111117bd5a3b060947b617355ecf353627c50b1xxxxxxxxxx
                   API Host: api-xxxxxxx.duosecurity.com
                   Autopush: true
                  Push info: true
                   Failmode: safe
                 Http-proxy: -
                    Prompts: 1
               Comment: This is a new key
```

security login duo group create

Add a Duo Group Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo group create` creates the groups that need to included in Duo authentication.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver of the Duo group that is being created.

-group-name <text> - Group Name

This parameter specifies the name of the group that will be part of Duo authentication.

[-excluded-users <text>,…] - List of Excluded Users

This parameter optionally specifies the list of users that will be excluded from Duo authentication.

[-comment <text>] - Comment

This parameter optionally specifies the comments

Examples

The following command creates a group "test" on vserver "VS1" and exclude the user tsmith from Duo authentication.

```
cluster1::> security login duo group create -vserver vs1 -group-name vs1
-exclude-users tsmith
```

security login duo group delete

Delete a Duo Group Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo group delete` command deletes the specified group.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver of the Duo group that is being deleted.

-group-name <text> - Group Name

This parameter optionally specifies the group name that is being deleted.

Examples

The following command deletes the group "test" from the Vserver "vs1".

```
cluster1::> security login duo group delete -vserver vs1 -group-name test
```

security login duo group modify

Modify a Duo Group Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo group modify` command modifies the Duo groups.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver of the group for which the 'exclude-users' or 'comment' is being modified.

-group-name <text> - Group Name

This parameter specifies the name of the group for which the 'exclude-users' or 'comment' is being modified.

[-excluded-users <text>,…] - List of Excluded Users

This parameter specifies the list of users to be excluded from Duo authentication.

[-comment <text>] - Comment

This parameter specifies the comments.

Examples

The following command modifies the excluded users for the group test who is part of Vserver vs1.

```
cluster1::> security login duo group modify -vserver vs1 -group-name test
-exclude-users tsmith, jane
```

security login duo group show

Display Duo Group Configurations

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login duo group show` command displays information about Duo groups.

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.

[-vserver <vserver name>] - Vserver

Selects the Vserver that match this parameter value.

[-group-name <text>] - Group Name

Selects the groups that match this parameter value.

[-excluded-users <text>,…] - List of Excluded Users

Selects the excluded users that match this parameter value.

[-comment <text>] - Comment

Selects the comments that match this parameter value.

Examples

The example below displays group information for the Vserver vs1.

```
cluster1::> security login duo group show -vserver vs1
Vserver: vs1
Group Name: NETAPP_ENG
Excluded Users: user1, user2, user12
Comment: This is a new group
```

security login external-role-mapping create

Add a external role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command creates a mapping between an external IDP (Identity Provider) role and an ONTAP role. This allows the system to authorize users based on their roles in the external IDP.

Parameters

-external-role <text> - External Role Name

The role name coming from the external IDP. This should match the role name as defined in the IDP, for which the external role mapping is being created. For example, 'Administrator'.

-provider {basic|keycloak|auth0|adfs|entra} - Provider Type

The name of the IDP provider. This should be the identifier for the IDP system. For example, 'entra'.

-ontap-role <text> - Ontap Role Name

The name of the ONTAP role that the external role should map to. This should be a role defined within ONTAP. For example, 'admin'.

[-comment <text>] - Comment

This parameter optionally specifies any comments or notes related to the external role mapping. This can be used to provide additional context or explanation for the mapping.

Examples

This command creates an external role mapping for the external role "Administrator" from the IDP provider "entra" to the ONTAP role "admin". The comment 'External role mapping' provides additional context for this mapping.

```
cluster1::> security login external-role-mapping create -external-role
Administrator -provider entra -ontap-role admin -comment "External role
mapping"
```

security login external-role-mapping delete

Delete a external role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command deletes an external role mapping.

Parameters

`-external-role <text>` - External Role Name

This parameter specifies the name of the external role to be deleted.

`-provider {basic|keycloak|auth0|adfs|entra}` - Provider Type

The IDP (Identity Provider) type of the external role to be deleted.

Examples

This command deletes an external role mapping named 'Administrator' from the IDP 'entra'.

```
cluster1::> security login external-role-mapping delete -external-role
Administrator -provider entra
```

security login external-role-mapping modify

Modify a external role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command modifies a configured mapping between an external IDP (Identity Provider) role and an ONTAP role.

Parameters

`-external-role <text>` - External Role Name

Identifies the external role mapping that is being modified.

`-provider {basic|keycloak|auth0|adfs|entra}` - Provider Type

Identifies the IDP provider associated with the external role mapping being modified.

[-ontap-role <text>] - Ontap Role Name

The new ONTAP role to which the external role should be mapped. This should be a role defined within ONTAP. For example 'admin'.

[-comment <text>] - Comment

Optionally, new comments or notes related to the external role mapping can be specified here.

Examples

This command modifies an external role mapping for the external role "Administrator" from the IDP provider "entra" to the ONTAP role "readonly". The comment 'Modified role mapping' provides additional context for this mapping.

```
cluster1::> security login external-role-mapping modify -external-role
Administrator -provider entra -ontap-role readonly -comment "Modified role
mapping"
```

security login external-role-mapping show

Display the external role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays the configured mapping between an external IDP (Identity Provider) role and an ONTAP role.

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.

[-external-role <text>] - External Role Name

Selects the external role mapping associated with the name that matches this parameter value.

[-provider {basic|keycloak|auth0|adfs|entra}] - Provider Type

Selects the external role mapping associated with the provider that matches this parameter value.

[-ontap-role <text>] - Ontap Role Name

Selects the external role mapping associated with the ONTAP role that matches this parameter value.

[-comment <text>] - Comment

Selects the external role mapping associated with the comment that matches this parameter value.

[-created <MM/DD/YYYY HH:MM:SS>] - Create Timestamp

Selects the external role mapping associated with the created time that matches this parameter value.

Examples

The example below displays information about the external role mapping for the role 'Administrator' from the IDP 'entra'.

```
cluster1::> security login external-rolemapping show -external-role
Testing -provider entra
External Role Name     Provider Type    Ontap Role Name       Comment
---------------------  ---------------  ---------------------- ------------
Administrator          entra            Custom_role           -
```

security login group create

Add a group configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command creates a user group for the login method. A group can only be created with the type `entra` .

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver with which the group is associated. If this parameter is specified, the group is associated with that specific Vserver. If not specified, the group is associated with the cluster.

-name <text> - Group Name

The name of the group.

-type <entra> - Group Type

This parameter specifies the IDP the group originates from.

[-uuid <UUID>] - Group UUID

Specifies the universally unique identifier of the group.

[-comment <text>] - Comment

This parameter optionally specifies any comments or notes related to the group.

Examples

This command creates a group named 'Eng_group' for the Vserver 'ontapcluster-1', with the specified UUID and a comment indicating it’s an 'Engineering Group'.

```
cluster1::> security login group create -vserver ontapcluster-1 -name
Eng_group -type entra -uuid 2ef26315-b95e-4778-9944-cf9e7b945d3d -comment
"Engineering Group"
```

security login group delete

Delete a group configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command deletes a group.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver associated with the group. If this parameter is specified, the group deletion applies to that specific Vserver only. If not specified, the cluster Vserver setting is used.

`-name <text>` - Group Name

The name of the group.

`-type <entra>` - Group Type

This parameter specifies the IDP the group originates from.

Examples

This command deletes a group named 'Eng_group' for the Vserver 'ontapcluster-1', with the type 'entra'.

```
cluster1::> security login group delete -name Eng_group -type entra
-vserver ontapcluster-1
```

security login group modify

Modify a group configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command modifies a user group for the login method.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver with which the group is associated.

-name <text> - Group Name

The name of the group.

-type <entra> - Group Type

This parameter specifies the IDP the group originates from.

[-uuid <UUID>] - Group UUID

Specifies the universally unique identifier of the group.

[-comment <text>] - Comment

This parameter optionally specifies any comments or notes related to the group.

Examples

This command modifies the UUID of the group 'Eng_group' on the Vserver 'ontapcluster-1'.

```
cluster1::> security login group modify -vserver ontapcluster-1 -name
Eng_group -type entra -uuid 2ef26315-1234-4778-9944-cf9e7b945d3d
```

security login group show

Display group configurations

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the configured groups.

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.

[-vserver <vserver name>] - Vserver

Selects the groups associated with the Vserver that matches this parameter value.

[-name <text>] - Group Name

Selects the groups associated with the name that matches this parameter value.

[-type <entra>] - Group Type

Selects the groups associated with the type that matches this parameter value.

[-id <integer>] - Group ID

Selects the groups associated with the group id that matches this parameter value.

[-uuid <UUID>] - Group UUID

Selects the groups associated with the uuid that matches this parameter value.

[-created <MM/DD/YYYY HH:MM:SS>] - Create Timestamp

Selects the groups associated with the created time that matches this parameter value.

[-comment <text>] - Comment

Selects the groups associated with the comment that matches this parameter value.

[-scope {svm|cluster}] - Cluster or SVM

Selects the groups associated with the scope that matches this parameter value.

Examples

The example below displays the "Eng_group" group information of the Vserver VS1.

```
cluster1::> security login group show -name Eng_group -vserver VS1
                    Vserver: VS1
                 Group Name: Eng_group
                   Group ID: 10
                 Group Type: entra
                    Comment: Engineering Group
```

security login group role-mapping create

Add a role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command creates a mapping between a group and a role.

Parameters

-group-id <integer> - Group ID

The group ID for which the role mapping is being created. The group ID can be obtained by executing the command 'security login group show'.

-role <text> - Role Name

The name of the role to which the group is being mapped.

[-comment <text>] - Comment

This parameter optionally specifies any comments or notes related to the group to role mapping.

Examples

This command creates a role mapping for the group ID "1" to the role "admin" and a comment indicating it's 'Group of admins'.

```
cluster1::> security login group role-mapping create -group-id 1 -role
admin -comment "Group of admins"
```

security login group role-mapping delete

Delete a role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command deletes a mapping between a group and a role.

Parameters

`-group-id <integer>` - Group ID

The group ID of the group-role mapping to delete.

`-role <text>` - Role Name

The role name of the group-role mapping to delete.

Examples

This command deletes a group to role mapping for the group-id 1 and the role "admin".

```
cluster1::> security login group role-mapping delete -group-id 1 -role
admin
```

security login group role-mapping modify

Modify a role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command creates a mapping between a group and a role.

Parameters

`-group-id <integer>` - Group ID

The group ID for which the role mapping is being modified. The group ID can be obtained by executing the command 'security login group show'.

-role <text> - Role Name

The name of the role to which the group is being mapped.

[-comment <text>] - Comment

This parameter optionally specifies any comments or notes related to the group to role mapping.

Examples

This command modifies the comment of the role mapping for the group ID "1" and role "admin" to a comment indicating it's 'Group of admins'.

```
cluster1::> security login group role-mapping modify -group-id 1 -role
admin -comment "Group of admins"
```

security login group role-mapping show

Display the role mapping

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays the configured group to role mappings.

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.

[-group-id <integer>] - Group ID

Selects the group role mapping associated with the group ID that matches this parameter value.

[-role <text>] - Role Name

Selects the group role mapping associated with the role name that matches this parameter value.

[-comment <text>] - Comment

Selects the group role mapping associated with the comment that matches this parameter value.

[-scope {svm|cluster}] - Cluster or SVM

Selects the group role mapping associated with the scope that matches this parameter value.

Examples

The example below displays the group to role mapping for the group ID 1 and the role "admin".

```
cluster1::> security login group role-mapping show -group-id 1 -role admin
                    Group ID: 1
                   Role Name: admin
                     Comment: -
              Cluster or SVM: cluster
```

security login motd modify

Modify the message of the day

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login motd modify` command updates the message of the day (MOTD).

There are two categories of MOTDs: the cluster-level MOTD and the data Vserver-level MOTD. A user logging in to a data Vserver's clustershell will potentially see two messages: the cluster-level MOTD followed by the Vserver-level MOTD for that Vserver. The cluster administrator can enable or disable the cluster-level MOTD on a per-Vserver basis. If the cluster administrator disables the cluster-level MOTD for a Vserver, a user logging into the Vserver will not see the cluster-level message. Only a cluster administrator can enable or disable the cluster-level message.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Use this parameter to specify the Vserver whose MOTD will be modified. Use the name of the cluster admin Vserver to modify the cluster-level message.

{ `[-message <text>]` - Message of the Day (MOTD)

This optional parameter can be used to specify a message. If you use this parameter, the MOTD cannot contain newlines (also known as end of lines (EOLs) or line breaks). If you do not specify any parameter other than the `-vserver` parameter, you will be prompted to enter the message interactively. Messages entered interactively can contain newlines. Non-ASCII characters must be provided as Unicode UTF-8.

The message may contain dynamically generated content using the following escape sequences:

- `\\` - A single backlash character.
- `\b` - No output: supported for compatibility with Linux only.
- `\C` - Cluster name.
- `\d` - Current date as set on the login node.
- `\t` - Current time as set on the login node.
- `\I` - Incoming LIF IP address (prints 'console' for a console login).
- `\l` - Login device name (prints 'console' for a console login).
- `\L` - Last login for the user on any node in the cluster.

- *\m* - Machine architecture.
- *\n* - Node or data Vserver name.
- *\N* - Name of user logging in.
- *\o* - Same as \O. Provided for Linux compatibility.
- *\O* - DNS domain name of the node. Note that the output is dependent on the network configuration and may be empty.
- *\r* - Software release number.
- *\s* - Operating system name.
- *\u* - Number of active clustershell sessions on the local node. For the cluster admin: all clustershell users. For the data Vserver admin: only active sessions for that data Vserver.
- *\U* - Same as \u, but has 'user' or 'users' appended.
- *\v* - Effective cluster version string.
- *\W* - Active sessions across the cluster for the user logging in ('who').

A backslash followed by any other character is emitted as entered.

| [-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}] - Download URI for the MOTD }

Use this parameter to specify the URI from where the message of the day will be downloaded. Note that the message must not exceed 2048 bytes in length. Non-ASCII characters must be provided as Unicode UTF-8.

[-is-cluster-message-enabled {true|false}] - Is Cluster-level Message Enabled?

Use this parameter to enable or disable the display of the cluster-level MOTD for the specified Vserver.

Examples

This example shows how to enter a MOTD interactively:

```
cluster1::> security login motd modify -vserver vs0

Enter the message of the day for Vserver "vs0".
Max size: 2048. Enter a blank line to terminate input. Press Ctrl-C to
abort.
0        1         2         3         4         5         6         7
8
12345678901234567890123456789012345678901234567890123456789012345678901234
567890
Welcome to the Vserver!
cluster1::>
```

security login motd show

Display the message of the day

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login motd show` command displays information about the cluster-level and data Vserver clustershell message of the day (MOTD).

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.

[-vserver <Vserver Name>] - Vserver Name

Selects the message of the day entries that match this parameter value. Use the name of the cluster admin Vserver to see the cluster-level MOTD.

[-message <text>] - Message of the Day (MOTD)

Selects the message of the day entries that match this parameter value.

[-is-cluster-message-enabled {true|false}] - Is Cluster-level Message Enabled?

Selects the message of the day entries that match this parameter value.

Examples

The following example displays all message of the day entries:

```
cluster1::> security login motd show
Vserver: cluster1
Is the Cluster MOTD Displayed?: true
Message
 ---------------------------------------------------------------------------
----
The cluster is running normally.

Vserver: vs0
Is the Cluster MOTD Displayed?: true
Message
 ---------------------------------------------------------------------------
----
Welcome to the Vserver!

2 entries were displayed.
```

security login publickey create

Add a new public key

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login publickey create` associates an existing public key with a user account. This command requires that you enter a valid OpenSSH-formatted public key, a user name, index number, and optionally, a comment and a certificate.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver of the user for whom you are adding the public key.

`-username <text>` - Username

This parameter specifies the name of the user for whom you are adding the public key. If you do not specify a user, the user named admin is specified by default.

`[-application <text>]` - Application

This parameter optionally specifies the application for which you are adding the public key. Presently, the only supported values are *`ssh`* and *`service-processor`* . The default value for this parameter is *`ssh`* .

`[-index <integer>]` - Index

This parameter specifies an index number for the public key. The default value is the next available index value, starting with zero if it is the first public key created for the user.

-publickey <certificate> - Public Key

This specifies the OpenSSH public key, which must be enclosed in double quotation marks.

[-comment <text>] - Comment

This optionally specifies comment text for the public key. Note that comment text should be enclosed in quotation marks.

[-x509-certificate <text>] - Install/Modify/Delete X509 Certificate

If this parameter is used, the specified certificate will be installed. The default when the public key is created is no certificate.

Examples

The following command associates a public key with a user named tsmith for Vserver vs1. The public key is assigned index number 5 and the comment text is “This is a new key”. Optionally, we can also specify the certificate.

```
cluster1::> security login publickey create -vserver vs1 -username tsmith
-index 5 -publickey
"ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAspH64CYbUsDQCdW22JnK6J
/vU9upnKzd2zAk9C1f7YaWRUAFNs2Qe5lUmQ3ldi8AD0Vfbr5T6HZPCixNAIza
FciDy7hgnmdj9eNGedGr/JNrftQbLD1hZybX+72DpQB0tYWBhe6eDJ1oPLob
ZBGfMlPXh8VjeU44i7W4+s0hG0E=tsmith@publickey.example.com"
-comment "This is a new key"
-x509-certificate install
Enter certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIBYTCB56ADAgECAhQybXS4lSL99APGRPqaGTy7CGUICzAKBggqhkjOPQQDAjAS
MRAwDgYDVQQDDAdTU0ggS2V5MB4XDTIyMDkyNjA1MDcwMVoXDTIzMDkyNjAwMDAw
MFowEjEQMA4GA1UEAwwHU1NIIEtleTB2MBAGByqGSM49AgEGBSuBBAAiA2IABEL4
UEJYUBfTO8gGSdRQleLQNvxVFjTiCN9V+8/6qsSTshb+K2zGT7qoF2RNPRMtHvtH
r/EC7Wo+9yP/ovtjFsshC+boJpfe8NN4xpqDzeC0nn1kw1GIavOCGyhGUNauITAK
BggqhkjOPQQDAgNpADBmAjEA4ClCSp3Nb7DlX6Bxvi7utQobj2qQETgNxHpiYzlD
Zr920lNPHDRxJTQ04vIdNeoZAjEAwTfv5jpuowaYxXxPJND2ytpyFcjyl/BUrFpQ
7XyjxyVFzKP3Rfj+uBvhIb8sLbl8
-----END CERTIFICATE-----
```

The following command associates a public key with a 'service-processor' user named 'joed' for the Administrative Vserver. The public key is assigned index number 1 and the comment text is "This is a new publickey". In this case, the certificate has not been specified.

```
cluster1::> security login publickey create -username joed -application
service-processor -index 1 -publickey
"ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBFgMRmKHV2+sGDPMD8JdMR
S+vSviIyjAymBZ/vu22Ae1AY51RIOBq5cVKaP49O9A21F1Srkqsd3gHFw/UhYBPy8="
-comment "This is a new publickey"
```

security login publickey delete

Delete a public key

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login publickey delete` command deletes a public key for a specific user. To delete a public key, you must specify a user name and index number.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver of the user for whom you are adding the public key.

`-username <text>` - Username

This parameter specifies the name of the user for whom you are deleting a public key. If you do not specify a user, the user named admin is specified by default.

`-application <text>` - Application

This parameter optionally specifies the application for which you are deleting the public key. Presently, the only supported values are *`ssh`* and *`service-processor`* . The default value for this parameter is *`ssh`* .

`-index <integer>` - Index

This parameter specifies an index number for the public key.

Examples

The following command deletes the public key for the user named tsmith with the index number 5 along with the certificate if it was installed.

```
cluster1::> security login publickey delete -username tsmith -application
ssh -index 5
```

The following command deletes the public key at the index number 2 for the 'service-processor' user named 'joed'.

```
cluster1::> security login publickey delete -vserver cluster1 -username
joed -application service-processor -index 2
```

security login publickey load-from-uri

Load one or more public keys from a URI

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login publickey load-from-uri` command loads one or more public keys from a Universal Resource Identifier (URI). To load public keys from a URI, you must specify a user name, the URI from which to load them, and optionally, whether you want to overwrite the existing public keys.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver for the user associated with the public keys.

`-username <text>` - Username

This parameter specifies the username for the public keys. If you do not specify a username, the username *"admin"* is used by default.

`-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - URI to load from

This parameter specifies the URI from which the public keys will be loaded.

`-overwrite {true|false}` - Overwrite Entries

This parameter optionally specifies whether you want to overwrite existing public keys. The default value for this parameter is *false* . If the value is *true* and you confirm to overwrite, then the existing public keys are overwritten with the new public keys. If you use the value *false* or do not confirm the overwrite, then newly loaded public keys are appended to the list of existing public keys using the next available index.

Examples

The following command shows how to load public keys for the user named tsmith from the URI ftp://ftp.example.com/identity.pub. This user's existing public keys are not overwritten.

```
cluster1::> security login publickey load-from-uri -username tsmith
  -uri ftp://ftp.example.com/identity.pub -overwrite false
```

The following command shows how to load public keys for the user named tsmith from the URI ftp:ftp://ftp.example.com/identity.pub. This user's existing public keys are overwritten if user entered the option 'y' or 'Y'. The user's existing public keys are not overwritten if user entered the option 'n' or 'N' and the newly loaded public keys are appended to the list of existing public keys using the next available index. The user and password credentials that you provide when you use this command are the credentials to access the server specified by the URI.

```
cluster1::> security login publickey load-from-uri -username
  tsmith -uri ftp://ftp.example.com/identity.pub -overwrite true -vserver
vs0

Enter User:
Enter Password:

Warning: You are about to overwrite the existing publickeys for the user
"tsmith" in Vserver "vs0". Do you want to proceed? {y|n}:
```

security login publickey modify

Modify a public key

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login publickey modify` command modifies a public key and optionally its comment text and certificate.

Parameters

`-vserver <vserver name>` - Vserver

Specifies the Vserver for the user associated with the public key.

`-username <text>` - Username

Specifies the username for the public key. If you do not specify a username, the username 'admin' is used by default.

`-application <text>` - Application

Optionally specifies the application for which you are modifying the public key. Presently, the only supported values are *`ssh`* and *`service-processor`* . The default value for this parameter is *`ssh`* .

`-index <integer>` - Index

Specifies the index number of the public key. The index number of the public key can be found by using the security login publickey show command.

`[-publickey <certificate>]` - Public Key

Specifies the new public key. You must enclose the new public key in double quotation marks.

`[-comment <text>]` - Comment

Specifies the new comment text for the public key.

`[-x509-certificate <text>]` - Install/Modify/Delete X509 Certificate

This parameter is used to modify or delete an existing certificate.

Examples

The following command modifies the public key and certificate at index number 10 for the SSH user named tsmith of Vserver vs1.

```
cluster1::> security login publickey modify -vserver vs1 -username tsmith
-index 10 -publickey
"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABAQDDD+pFzFgV/2dlowKRFgym9K910H/u+BVTGitCtHteHy
o8thmaXT
1GLCzaoC/12+XXiYKMRhJ00S9Svo4QQKUXHdCPXFSgR5PnAs39set39ECCLzmduplJnkWtX96p
QH/bg2g3upFcdC6z9
c37uqFtNVPfv8As1Si/9WDQmEJ2mRtJudJeU5GZwZw5ybgTaN1jxDWus9SO2C43F/vmoCKVT52
9UHt4/ePcaaHOGTiQ
O8+Qmm59uTgcfnpg53zYkpeAQV8RdYtMdWlRr44neh1WZrmW7x5N4nXNvtEzr9cvb9sJyqTX1C
kQGfDOdb+7T7y3X7M
if/qKQY6FsovjvfZD" -x509-certificate modify
Enter certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIBYjCB6aADAgECAhRhhyz6AuV04M5nac2Pq+cdm4IRMzAKBggqhkjOPQQDAjAT
MREwDwYDVQQDDAhhYmhpamVlMjAeFw0yMjEwMTkxNzI4MzZaFw0yMjEwMjAwMDAw
MDBaMBMxETAPBgNVBAMMCGFiaGlqZWUyMHYwEAYHKoZIzj0CAQYFK4EEACIDYgAE
xNEN5jC3IjAvR0kwsAyKXw7aBLwNg5nvDC1D2bRLpbZkoLQ6hzyKZUEqA0ELlQla
u4yX9O1qiRIbMMclqMB1XYJywxOh+3uKOTM6Bz82/IZp4Oaa/4gYVtFRgStHTdPf
MAoGCCqGSM49BAMCA2gAMGUCMQDQIgdSrECuWJ76ZvfEDAvFlHnJHQtNz8zFl0lh
XRNzPlhpltEm6j5V6mPkRmJrmloCMGckXAVkmUCGFxU2e2ZvuKbL5BVCrE5iifet
ly3UApgBg8EgTO+hebMNz3i/Z4p+5w==
-----END CERTIFICATE-----"
```

The following command modifies the public key at index number 3 for the 'service-processor' user named 'joed' belonging to the Administrative Vserver.

```
cluster1::> security login publickey modify -vserver cluster1 -username
joed -index 3 -publickey
"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABAQDYaIWUyYFtW+8xeCFgVg5OEM2P4m6mZvOBj7ES6TTUSI
cYk4gaOyLcJNwm05Q+Op3tidvakx16ALCtkP9TBBWgHPwwHma0+z689ekK88myabMst12iwMR3
9OGmen2KKW7NJQwwWmed47tTkEi5VU1BfX3CAsXiw4jp1HuFYi2iuHUQDtN6MlL4ON51X2IYUv
StH4N1UNn89bn3Q0l4UYFpFwF2ixIuR8Pm11bZvDy1yP4hpxmoisUpkwsmr/SJiBsX381Ogg3C
+Gnthtqe5/xoYx2CCb93Ff2UqKQyf41MgK8P1gNpOt5Vns1LQ+K1S+fh4ZmtkmBwHvQEXwBKQg
mb"
```

Related Links

- security login publickey show

security login publickey show

Display public keys

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login publickey show` command displays information about public keys.

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.

[-vserver <vserver name>] - Vserver

Selects the public keys that match this parameter value.

[-username <text>] - Username

Selects the public keys that match this parameter value.

[-application <text>] - Application

Selects the public keys that match this parameter value.

[-index <integer>] - Index

Selects the public keys that match this parameter value.

[-publickey <certificate>] - Public Key

Selects the public keys that match this parameter value.

[-fingerprint <text>] - Hex Fingerprint

Selects the public keys that match this parameter value.

[-bubblebabble <text>] - Bubblebabble Fingerprint

Selects the public keys that match this parameter value.

[-comment <text>] - Comment

Selects the public keys that match this parameter value.

[-certificate <certificate>] - Certificate Associated with Public Key

Selects the public keys that match this parameter value.

[-certificate-details <text>] - Details about the Certificate

Selects the public keys that match this parameter value.

`[-certificate-expired <text>]` - Expiry Status of Certificate

Selects the public keys that match this parameter value.

`[-certificate-revoked <text>]` - Revocation Status of Certificate

Selects the public keys that match this parameter value.

Examples

The example below displays public key information for the user named tsmith.

```
cluster1::> security login publickey show -username tsmith
Vserver: vs1
UserName: tsmith Index: 5
Public Key:
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAspH64CYbUsDQCdW22JnK6J
/vU9upnKzd2zAk9C1f7YaWRUAFNs2Qe5lUmQ3ldi8AD0Vfbr5T6HZPCixNAIza
FciDy7hgnmdj9eNGedGr/JNrftQbLD1hZybX+72DpQB0tYWBhe6eDJ1oPLob
ZBGfMlPXh8VjeU44i7W4+s0hG0E=tsmith@publickey.example.com
Fingerprint:
07:b4:27:52:ce:7f:35:81:5a:f2:07:cf:c1:87:91:97
Bubblebabble fingerprint:
xuzom-nelug-bisih-nihyr-metig-kemal-puhut-somyd-mumuh-zomis-syxex
Comment:
This is a new key
Certificate:
-----BEGIN CERTIFICATE-----
MIICHTCCAaKgAwIBAgIUU+YJjeaOzvs+w+56JSm8Amow0nAwCgYIKoZIzj0EAwIw
RTELMAkGA1UEBhMCSU4xCzAJBgNVBAgMAktBMQwwCgYDVQQHDANCTFIxDTALBgNV
BAoMBE5UQVAxDDAKBgNVBAsMA1ImRDAeFw0yMjA5MjAwNjAyNDRaFw0yMzA5MjAw
NjAyNDRaMEUxCzAJBgNVBAYTAklOMQswCQYDVQQIDAJLQTEMMAoGA1UEBwwDQkxS
MQ0wCwYDVQQKDAROVEFQMQwwCgYDVQQLDANSJkQwdjAQBgcqhkjOPQIBBgUrgQQA
IgNiAASbaxCYB6XRDOFgdHBghvMoUpideGnd2jNrQJANeSaWVMnPUpXzg2tcPnsu
c87AR75BcfwhSurrFGLIw7TLcR22IFTggcrKmhjI8QwvomMZWfIoeHZlwsI+msm4
PLyaCqmjUzBRMB0GA1UdDgQWBBREeGdLZ3YTEL4CXLvTa8XQRAhRqDAfBgNVHSME
GDAWgBREeGdLZ3YTEL4CXLvTa8XQRAhRqDAPBgNVHRMBAf8EBTADAQH/MAoGCCqG
SM49BAMCA2kAMGYCMQDv9ZEselgteBlbqOYScKCyVcq3d89zz8Y9GBBB4FXJ3J+q
/h4zDk2Y2IJG63d7Kf0CMQDJl7v9I/NRNtS09qkavJh6snjJvUe3C5RhAkMPDhBO
2sfbUx1UQSo/md6U1CQBewM=
-----END CERTIFICATE-----
Certificate Details:
Subject:
C=US, ST=NC, L=RTP, O=NETAPP, OU=NTAP,
CN=scspr2692789021.gdl.englab.netapp.com
Issuer:
C=US, ST=NC, L=RTP, O=NETAPP, OU=NTAP, CN=NTAP-INTERCA2
Expiration: Jan 29 04:46:20 2024 GMT
Certificate Expiration Status: Not Expired
Certificate Revocation Status: good
```

The example below displays public key information for all 'service-processor' users.

```
cluster1::> security login publickey show -application service-processor
Vserver: cluster1
UserName: joed Index: 1
Public Key:
ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBFgMRmKHV2+sGDPMD8JdMR
S+vSviIyjAymBZ/vu22Ae1AY51RIOBq5cVKaP49O9A21F1Srkqsd3gHFw/UhYBPy8=
Fingerprint:
SHA256:ETxoqJyjq2tXjMKiJlCSTOT5vj+s+h6OxtXUIl28PcI
Bubblebabble fingerprint:
xiges-husyn-fyzim-sanok-bihos-sizuv-ribyt-cyryz-lelel-sekan-poxyx
Comment:
This is a new publickey
Certificate:
-
Certificate Details:
-
Certificate Expiration Status: -
Certificate Revocation Status: -
```

security login rest-role create

Add a REST access control role

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login rest-role create` command creates a Representational State Transfer (REST) access-control role. A REST access-control role consists of a role name and an Application Programming Interface (API) to which the role has access. It optionally includes an access level (*`none`* , *`readonly`* , *`read_create`* , *`read_modify`* , *`read_create_modify`* or *`all`*) for the API. After you create a REST access-control role, you can apply it to a management-utility login account by using the security login modify or security login create commands.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name associated with the REST role.

`-role <text>` - Role Name

This specifies the REST role that is to be created.

`-api <text>` - API Path

This specifies the API to which the REST role has access. This API can be a private CLI API or a resource-qualified endpoint. Currently, the only supported resource-qualified endpoints are the following:

- Snapshots APIs
- */api/storage/volumes/{volume.uuid}/snapshots*
- File System Analytics APIs
- */api/storage/volumes/{volume.uuid}/files*
- */api/storage/volumes/{volume.uuid}/top-metrics/clients*
- */api/storage/volumes/{volume.uuid}/top-metrics/directories*
- */api/storage/volumes/{volume.uuid}/top-metrics/files*
- */api/storage/volumes/{volume.uuid}/top-metrics/users*
- */api/svm/svms/{svm.uuid}/top-metrics/clients*
- */api/svm/svms/{svm.uuid}/top-metrics/directories*
- */api/svm/svms/{svm.uuid}/top-metrics/files*
- */api/svm/svms/{svm.uuid}/top-metrics/users*
- ONTAP S3 APIs
- */api/protocols/s3/services/{svm.uuid}/users*

In the above APIs, wildcard character *** could be used in place of *{volume.uuid}* or *{svm.uuid}* to denote *all* volumes or *all* SVMs, depending upon whether the REST endpoint references volumes or SVMs.

-access {none|readonly|read_create|read_modify|read_create_modify|all} - Access Level

This optionally specifies an access level for the REST role. Possible access level settings are *none* , *readonly* , *read_create* , *read_modify* , *read_create_modify* and *all* .

Examples

The following command creates a REST access-control role named *admin* for the *vs1.example.com* Vserver. This REST role has an access-level of *all* for the */api/storage/volumes* API.

```
cluster1::> security login rest-role create -role admin -api
"/api/storage/volumes" -access all -vserver vs1.example.com
cluster1::>
```

The following command creates a REST access-control role named *rest_role1* for the *cluster1.example.com* administrative Vserver. This REST role has an access-level of *read_create_modify* for the */api/snapmirror/policies* API.

```
cluster1::> security login rest-role create -role rest_role1 -api
"/api/snapmirror/policies" -access read_create_modify -vserver
cluster1.example.com
cluster1::>
```

The following command creates a REST access-control role named *vs1_role* for the *vs1.example.com* Vserver. This REST role has an access level of *readonly* for all snapshots on the volume with UUID *f8a541b5-b68c-11ea-9581-005056bbabe6* . The volume UUID refers to the `-instance-uuid` field value in the volume show command output at diagnostic privilege level.

```
cluster1::> security login rest-role create -role vs1_role -api
"/api/storage/volumes/f8a541b5-b68c-11ea-9581-005056bbabe6/snapshots"
-access readonly -vserver vs1.example.com
Warning: Operating on an alias operates on the target of the specified
alias:
           "volume snapshot"
cluster1::>
```

The following command creates a REST access-control role named *vs2_role* for the *vs2.example.com* Vserver. This REST role has an access level of *readonly* for all files on the volume with UUID *15d489b5-1d40-11ec-992e-005056bba268* .The volume UUID refers to the `-instance-uuid` field value in the volume show command output at diagnostic privilege level.

```
cluster1::> security login rest-role create -role vs2_role -api
"/api/storage/volumes/15d489b5-1d40-11ec-992e-005056bba268/files" -access
readonly -vserver vs2.example.com
cluster1::>
```

The following command creates a REST access-control role named *vs3_role* for the *vs3.example.com* Vserver. This REST role has an access level of *read_create_modify* for all top-metrics directories on the SVM with UUID *881764b5-9ea1-11ec-8771-005056bb1a7c* .

```
cluster1::> security login rest-role create -role vs3_role -api
"/api/svm/svms/881764b5-9ea1-11ec-8771-005056bb1a7c/top-
metrics/directories" -access read_create_modify -vserver vs3.example.com
cluster1::>
```

The following command creates a REST access-control role named *vs4_role* for the *vs4.example.com* Vserver. This REST role has an access level of *all* for command directory *cluster* .

```
cluster1::> security login rest-role create -role vs4_role -api
"/api/private/cli/cluster" -access all -vserver vs4.example.com
cluster1::>
```

Related Links

- security login modify
- security login create

- volume show

security login rest-role delete

Delete a REST access control role

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login rest-role delete` command deletes a Representational State Transfer (REST) access-control role.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name associated with the REST role.

`-role <text>` - Role Name

This specifies the REST role that is to be deleted.

`-api <text>` - API Path

This specifies the Application Programming Interface (API) to which the REST role has access. This API can be a private CLI API or a resource-qualified endpoint. Currently, the only supported resource-qualified endpoints are the following:

- Snapshots APIs
- *`/api/storage/volumes/{volume.uuid}/snapshots`*
- File System Analytics APIs
- *`/api/storage/volumes/{volume.uuid}/files`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/clients`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/directories`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/files`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/users`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/clients`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/directories`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/files`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/users`*
- ONTAP S3 APIs
- *`/api/protocols/s3/services/{svm.uuid}/users`*

In the above APIs, wildcard character `*` could be used in place of *`{volume.uuid}`* or *`{svm.uuid}`* to denote *`all`* volumes or `_all _` SVMs, depending upon whether the REST endpoint references volumes or SVMs.

Examples

The following command deletes a REST access-control role entry with the role name *readonly* and the API */api/storage/volumes* from Vserver *vs.example.com*.

```
cluster1::> security login rest-role delete -role readonly -api
"/api/storage/volumes" -vserver vs.example.com
cluster1::>
```

The following command deletes a REST access-control role entry with the role name *vs1_role* and the resource-qualified endpoint corresponding to all snapshots on the volume with UUID *0aa39ec1-b68d-11ea-9581-005056bbabe6* from Vserver *vs1.example.com*.The volume UUID refers to the `-instance-uuid` field value in the volume show command output at diagnostic privilege level.

```
cluster1::> security login rest-role delete -role vs1_role -api
"/api/storage/volumes/0aa39ec1-b68d-11ea-9581-005056bbabe6/snapshots"
-vserver vs1.example.com
cluster1::>
```

The following command deletes a REST access-control role entry with the role name *vs2_role* and the resource-qualified endpoint corresponding to all top-metrics clients on the volume with UUID *373eb9ef-1d40-11ec-992e-005056bba268* from Vserver *vs2.example.com*. The volume UUID refers to the `-instance-uuid` field value in the volume show command output at diagnostic privilege level.

```
cluster1::> security login rest-role delete -role vs2_role -api
"/api/storage/volumes/373eb9ef-1d40-11ec-992e-005056bba268/top-
metrics/clients" -vserver vs2.example.com
cluster1::>
```

The following command deletes a REST access-control role entry with the role name *vs3_role* and the resource-qualified endpoint corresponding to all top-metrics directories for the Vserver *vs3.example.com* with UUID *6dfeb2a7-9a16-11ec-819e-005056bb1a7c*.

```
cluster1::> security login rest-role delete -role vs3_role -api
"/api/svm/svms/6dfeb2a7-9a16-11ec-819e-005056bb1a7c/top-
metrics/directories" -vserver vs3.example.com
cluster1::>
```

The following command deletes a REST access-control role entry with the role name *vs4_role* and the API */api/private/cli/cluster* for the Vserver *vs4.example.com*.

```
cluster1::> security login rest-role delete -role vs4_role -api
"/api/private/cli/cluster" -vserver vs4.example.com
cluster1::>
```

Related Links

- volume show

security login rest-role modify

Modify a REST access control role

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login rest-role modify` command modifies a Representational State Transfer (REST) access-control role.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name associated with the REST role.

`-role <text>` - Role Name

This specifies the REST role that is to be modified.

`-api <text>` - API Path

This specifies the Application Programming Interface (API) to which the REST role has access. This API can be a private CLI API or a resource-qualified endpoint. Currently, the only supported resource-qualified endpoints are the following:

- Snapshots APIs
- *`/api/storage/volumes/{volume.uuid}/snapshots`*
- File System Analytics APIs
- *`/api/storage/volumes/{volume.uuid}/files`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/clients`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/directories`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/files`*
- *`/api/storage/volumes/{volume.uuid}/top-metrics/users`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/clients`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/directories`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/files`*
- *`/api/svm/svms/{svm.uuid}/top-metrics/users`*

- ONTAP S3 APIs
- */api/protocols/s3/services/{svm.uuid}/users*

In the above APIs, wildcard character *** could be used in place of *{volume.uuid}* or *{svm.uuid}* to denote *all* volumes or *all* SVMs, depending upon whether the REST endpoint references volumes or SVMs.

[-access {none|readonly|read_create|read_modify|read_create_modify|all}] - Access Level

This specifies a new access level for the REST role. Possible access level settings are *none* , *readonly* , *read_create* , *read_modify* , *read_create_modify* and *all* .

Examples

The following command modifies a REST access-control role with the role name *readonly* and the API */api/storage/volumes* to have the access level *readonly* for Vserver *vs.example.com* :

```
cluster1::> security login rest-role modify -role readonly -api
"/api/storage/volumes" -access readonly -vserver vs.example.com
cluster1::>
```

The following command modifies a REST access-control role with the role name *rest_role1* and the API */api/snapmirror/policies* to have the access level *read_create* for Vserver *cluster1.example.com* :

```
cluster1::> security login rest-role modify -role rest_role1 -api
"/api/snapmirror/policies" -access read_create -vserver
cluster1.example.com
cluster1::>
```

The following command modifies a REST access-control role with the role name *vs1_role* and the resource-qualified endpoint */api/storage/volumes/*/snapshots* to have the access level *readonly* for Vserver *vs1.example.com* :

```
cluster1::> security login rest-role modify -role vs1_role -api
"/api/storage/volumes/*/snapshots" -access readonly -vserver
vs1.example.com
cluster1::>
```

The following command modifies a REST access-control role with the role name *vs2_role* and the resource-qualified endpoint */api/storage/volumes/4d383f47-1d40-11ec-81af-005056bb3eae/top-metrics/users* to have the access level *none* for Vserver *vs2.example.com* :

```
cluster1::> security login rest-role modify -role vs2_role -api
"/api/storage/volumes/4d383f47-1d40-11ec-81af-005056bb3eae/top-
metrics/users" -access none -vserver vs2.example.com
cluster1::>
```

The following command modifies a REST access-control role with the role name *vs3_role* and the resource-qualified endpoint */api/svm/svms/6dfeb406-9a16-11ec-819e-005056bb1a7c/top-metrics/files* to have the access level *read_modify* for Vserver *vs3.example.com*:

```
cluster1::> security login rest-role modify -role vs3_role -api
"/api/svm/svms/6dfeb406-9a16-11ec-819e-005056bb1a7c/top-metrics/files"
-access read_modify -vserver vs3.example.com
cluster1::>
```

The following command modifies a REST access-control role with the role name *vs3_role2* and the wildcard resource-qualified endpoint */api/svm/svms/*/top-metrics/clients* to have the access level *readonly* for Vserver *vs3.example.com*:

```
cluster1::> security login rest-role modify -role vs3_role2 -api
"/api/svm/svms/*/top-metrics/clients" -access readonly -vserver
vs3.example.com
cluster1::>
```

The following command modifies a REST access-control role with the role name *vs4_role* and api */api/private/cli/cluster* to have the access level *readonly* for Vserver *vs4.example.com*:

```
cluster1::> security login rest-role modify -role vs4_role -api
"/api/private/cli/cluster" -access readonly -vserver vs4.example.com
cluster1::>
```

security login rest-role show

Show REST access control roles

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login rest-role show` command displays the following information about Representational State Transfer (REST) access-control roles:

- Vserver
- Role name

- Application Programming Interface (API) to which the REST role has access
- Access Level (*none* , *readonly* , *read_create* , *read_modify* , *read_create_modify* , or *all*)

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.

[-vserver <vserver name>] - Vserver

Selects the REST roles that match this parameter value.

[-role <text>] - Role Name

Selects the REST roles that match this parameter value. If this parameter and the `-api` parameter are both used, the command displays detailed information about the specified REST access-control role.

[-api <text>] - API Path

Selects the REST roles that match this parameter value. If this parameter and the `-role` parameter are both used, the command displays detailed information about the specified REST access-control role. This API can be a private CLI API or a resource-qualified endpoint. Currently, the only supported resource-qualified endpoints are the following:

- Snapshots APIs
- */api/storage/volumes/{volume.uuid}/snapshots*
- File System Analytics APIs
- */api/storage/volumes/{volume.uuid}/files*
- */api/storage/volumes/{volume.uuid}/top-metrics/clients*
- */api/storage/volumes/{volume.uuid}/top-metrics/directories*
- */api/storage/volumes/{volume.uuid}/top-metrics/files*
- */api/storage/volumes/{volume.uuid}/top-metrics/users*
- */api/svm/svms/{svm.uuid}/top-metrics/clients*
- */api/svm/svms/{svm.uuid}/top-metrics/directories*
- */api/svm/svms/{svm.uuid}/top-metrics/files*
- */api/svm/svms/{svm.uuid}/top-metrics/users*
- ONTAP S3 APIs
- */api/protocols/s3/services/{svm.uuid}/users*
- Private-cli APIs
- */api/private/cli/cluster*

In the above APIs, wildcard character `*` could be used in place of *{volume.uuid}* or *{svm.uuid}* to

denote *all* volumes or *all* SVMs, depending upon whether the REST endpoint references volumes or SVMs.

[-access {none|readonly|read_create|read_modify|read_create_modify|all}] - Access Level

Selects the roles that match this parameter value.

Examples

The example below displays information about all REST access-control roles:

```
cluster1::> security login rest-role show
Role                             Access
      Vserver           Name            API             Level
      ----------        -------------   -----------     -----------
      vs                vsrole1         /api            none
      vs                vsrole1         /api/storage/volumes/f8a541b5-
b68c-11ea-9581-005056bbabe6/files
                                                        all
      vs                vsrole1         /api/storage/volumes/f8a541b5-
b68c-11ea-9581-005056bbabe6/snapshots
                                                        readonly
      vs                vsrole1         /api/storage/volumes/843b87f9-
2f5e-11ec-9524-005056bb0bee/snapshots
                                                        read_create
      vs                vsrole1         /api/svm/svms/843b87f9-2f5e-11ec-
9524-005056bb0bee/top-metrics/clients
                                                        read_create
      cluster1          readonly        /api/storage    none
      cluster1          custom          /api/cluster    read_modify
      cluster1          custom          /api/security/accounts
                                                        read_create_modify
      cluster1          custom          /api/storage/volumes/*/top-
metrics/users
                                                        readonly
      cluster1          custom          /api/storage/volumes/*/snapshots
                                                        all
cluster1::>
```

security login rest-role expanded-rest-roles modify

Modify the status of Expanded REST roles for granular resource control feature

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security login rest-role expanded-rest-roles modify` command enables or disables *`Expanded REST roles for granular resource control`* feature.

Parameters

`[-is-enabled {true|false}]` - Is Enabled? (privilege: advanced)

This parameter specifies whether the *`Expanded REST roles for granular resource control`* feature is enabled or disabled. The default value is *`true`* i.e. the feature is enabled by default.

Examples

The following command disables the *`Expanded REST roles for granular resource control`* feature.

```
cluster1::*> security login rest-role expanded-rest-roles modify -is
-enabled false
cluster1::*>
```

security login rest-role expanded-rest-roles show

Show the status of Expanded REST roles for granular resource control feature

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security login rest-role expanded-rest-roles show` command specifies whether the *`Expanded REST roles for granular resource control`* feature is enabled (*`true`*) or disabled (*`false`*).

Examples

The command below specifies that the *`Expanded REST roles for granular resource control`* feature is enabled.

```
cluster1::> security login rest-role expanded-rest-roles show

Is Enabled? true
```

security login role create

Add an access control role

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login role create` command creates an access-control role. An access-control role consists of a role name and a command or directory to which the role has access. It optionally includes an access level (none, readonly, or all) and a query that applies to the specified command or command directory. After you create an access-control role, you can apply it to a management-utility login account by using the security login modify or security login create commands.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name associated with the role.

`-role <text>` - Role Name

This specifies the role that is to be created.

`-cmddirname <text>` - Command / Directory

This specifies the command or command directory to which the role has access. The command or command directory must be specified either within double quotes or inside curly brackets. To specify the default setting, use the special value "`DEFAULT` ".

`[-access {none|readonly|read_create|read_modify|read_create_modify|all}]` - Access Level

This optionally specifies an access level for the role. Possible access level settings are none, readonly, and all. The default setting is `all` .

`[-query <query>]` - Query

This optionally specifies the object that the role is allowed to access. The query object must be applicable to the command or directory name specified by -cmddirname. The query object must be enclosed in double quotation marks (""), and it must be a valid field name.

Examples

The following command creates an access-control role named "admin" for the vs1.example.com Vserver. The role has all access to the "volume" command but only within the "aggr0" aggregate.

```
cluster1::> security login role create -role admin -cmddirname volume
-query "-aggr aggr0" -access all -vserver vs1.example.com
```

Related Links

- security login modify
- security login create

security login role delete

Delete an access control role

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login role delete` command deletes an access-control role.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name associated with the role.

`-role <text>` - Role Name

This specifies the role that is to be deleted.

`-cmddirname <text>` - Command / Directory

This specifies the command or command directory to which the role has access. To specify the default setting, use the special value "`DEFAULT` ".

Examples

The following command deletes an access-control role with the role name readonly and the command access "volume" for Vserver vs.example.com.

```
cluster1::> security login role delete -role readonly -cmddirname volume
-vserver vs.example.com
```

security login role modify

Modify an access control role

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login role modify` command modifies an access-control role.

Parameters

`-vserver <vserver name>` - Vserver

This optionally specifies the Vserver name associated with the role.

`-role <text>` - Role Name

This specifies the role that is to be modified.

`-cmddirname <text>` - Command / Directory

This specifies the command or command directory to which the role has access. To specify the default setting for a role, use the special value `"DEFAULT"` . This value can be modified only for the roles created for the admin Vserver.

`[-access {none|readonly|read_create|read_modify|read_create_modify|all}]` - Access Level

This optionally specifies a new access level for the role. Possible access level settings are none, readonly,

and all. The default setting is `all` .

[-query <query>] - Query

This optionally specifies the object that the role is allowed to access. The query object must be applicable to the command or directory name specified by -cmddirname. The query object must be enclosed in double quotation marks (""), and it must be a valid field name.

Examples

The following command modifies an access-control role with the role name readonly and the command access "volume" to have the access level readonly for Vserver vs.example.com:

```
cluster1::> security login role modify -role readonly -cmddirname volume
-access readonly -vserver vs.example.com
```

security login role prepare-to-downgrade

Update role configurations so that they are compatible with earlier releases of Data ONTAP

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security login role prepare-to-downgrade` command restores predefined roles of all Vservers earlier than ONTAP 8.3.2. You must run this command in advanced privilege mode when prompted to do so during the release downgrade.

Examples

The following command restores predefined roles of all Vservers earlier than ONTAP 8.3.2.

```
cluster1::*> security login role prepare-to-downgrade
```

security login role show-ontapi

Display the mapping between Data ONTAP APIs and CLI commands

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login role show-ontapi` command displays ONTAP APIs (ONTAPIs) and the CLI commands that they are mapped to.

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.

[-ontapi <text>] - ONTAPI Name

Use this parameter to view the corresponding CLI command for the specified API.

[-command <text>] - CLI Command

Use this parameter to view the corresponding API or APIs for the specified CLI command.

Examples

The following command displays all ONTAP APIs and their mapped CLI commands:

```
cluster1::> security login role show-ontapi
ONTAPI                      Command
---------------------------
-----------------------------------------------------
aggr-add                    storage aggregate add-disks
aggr-check-spare-low        storage aggregate check_spare_low
aggr-create                 storage aggregate create
aggr-destroy                storage aggregate delete
aggr-get-filer-info         aggr
aggr-get-iter               storage aggregate show-view
aggr-offline                storage aggregate offline
aggr-online                 storage aggregate online
aggr-options-list-info      storage aggregate show
aggr-rename                 storage aggregate rename
aggr-restrict               storage aggregate restrict
aggr-set-option             storage aggregate modify
autosupport-budget-get      system node autosupport budget show
autosupport-budget-get-iter system node autosupport budget show
autosupport-budget-get-total-records
                            system node autosupport budget show
autosupport-budget-modify   system node autosupport budget modify
autosupport-config-get      system node autosupport show
autosupport-config-get-iter system node autosupport show
autosupport-config-get-total-records
                            system node autosupport show
autosupport-config-modify   system node autosupport modify
Press <space> to page down, <return> for next line, or 'q' to quit...
```

The following example displays all ONTAP APIs which are mapped to the specified CLI command:

```
cluster1::> security login role show-ontapi -command version
ONTAPI                      Command
---------------------------
---------------------------------------------------
system-get-ontapi-version   version
system-get-version          version
2 entries were displayed.
```

The following example displays the CLI command that is mapped to the specified Data ONTAPI API:

```
cluster1::> security login role show-ontapi -ontapi aggr-create

ONTAPI Name: aggr-create
    Command: storage aggregate create
```

security login role show-rest

Display the mapping between ONTAP REST APIs and CLI commands

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login role show-rest` command displays ONTAP REST APIs and the CLI commands that they are mapped to.

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.

[-endpoint <text>] - URI of the REST endpoint

This parameter will filter the output to show CLI commands that match the provided REST endpoint.

[-commands <text>,…] - CLI Command

This parameter will filter the output to show the REST API or APIs that match the provided CLI command.

[-supported-operations <text>,…] - Supported REST operations <GET,POST,PATCH,DELETE>

This parameter filters the output to show REST APIs or CLI commands that match the provided REST operations <GET,POST,PATCH,DELETE>

Examples

The following command displays all the ONTAP REST APIs and their mapped CLI commands:

```
cluster1::> security login role show-rest
REST Endpoint               CLI Commands                    Supported
Operations
--------------------------- ------------------------------
--------------------
/application/applications   application                     GET, PATCH,
POST, DELETE
/application/applications/$id/components/$id/metadata
                            -                               GET, PATCH,
POST, DELETE
/application/applications/$id/metadata
                            application metadata            GET, PATCH,
POST, DELETE
/application/applications/{application.uuid}/components
                            -                               GET
/application/applications/{application.uuid}/components/{component.uuid}/s
napshots
                            -                               GET, POST,
DELETE
/application/applications/{application.uuid}/components/{component.uuid}/s
napshots/{uuid}/restore
                            -                               POST, DELETE
/application/applications/{application.uuid}/snapshots
                            -                               GET, POST,
DELETE
/application/applications/{application.uuid}/snapshots/{uuid}/restore
                            -                               POST, DELETE
/application/consistency-groups
Press <space> to page down, <return> for next line, or 'q' to quit...
```

The following examples display all ONTAP REST APIs that are mapped to a specified CLI command:

```
cluster1::> security login role show-rest -commands statistics
REST Endpoint               CLI Commands                    Supported
Operations
--------------------------- ------------------------------
--------------------
/cluster/counter/tables     statistics                      GET
/cluster/counter/tables/{counter_table.name}/rows
                            statistics                      GET
2 entries were displayed.
```

```
cluster1::> security login role show-rest -commands "network interface"*
REST Endpoint               CLI Commands                    Supported
Operations
--------------------------- ------------------------------
--------------------
/cluster/nodes              cluster add-node, network interface create,
storage aggregate auto-provision, system node, system node modify, system
node show, system service-processor network modify
                                                            GET, PATCH,
POST, DELETE
/network/fc/interfaces      network interface create, network interface
show, vserver fcp interface show
                                                            GET, PATCH,
POST, DELETE
/network/fc/interfaces/{uuid}/metrics
                            network interface show          GET
/network/ip/interfaces      network interface, network interface create
                                                            GET, PATCH,
POST, DELETE
/network/ip/interfaces/{uuid}/metrics
                            network interface show          GET
/network/ip/service-policies
                            network interface service-policy, network
interface service-policy show
                                                            GET, PATCH,
POST, DELETE
/svm/svms                   network interface create, network route
create, vserver, vserver add-aggregates, vserver add-protocols, vserver
cifs create, vserver create, vserver fcp create, vserver iscsi create,
vserver nfs create, vserver nvme create, vserver object-store-server
create, vserver remove-protocols, vserver services name-service dns
create, vserver services name-service ldap client create, vserver services
name-service nis-domain create, vserver services name-service ns-switch
create, vserver show
                                                            GET, PATCH,
POST, DELETE
7 entries were displayed.
```

```
cluster1::> security login role show-rest -commands *"fcp"*"create"
REST Endpoint               CLI Commands                   Supported
Operations
--------------------------- ------------------------------
--------------------
/protocols/san/fcp/services vserver fcp create, vserver fcp show
                                                           GET, PATCH,
POST, DELETE
/svm/svms                   network interface create, network route
create, vserver, vserver add-aggregates, vserver add-protocols, vserver
cifs create, vserver create, vserver fcp create, vserver iscsi create,
vserver nfs create, vserver nvme create, vserver object-store-server
create, vserver remove-protocols, vserver services name-service dns
create, vserver services name-service ldap client create, vserver services
name-service nis-domain create, vserver services name-service ns-switch
create, vserver show
                                                           GET, PATCH,
POST, DELETE
2 entries were displayed.
```

The following examples display the CLI command that is mapped to a specified ONTAP REST API:

```
cluster1::> security login role show-rest -endpoint /cluster/metrocluster*
REST Endpoint               CLI Commands                    Supported
Operations
--------------------------- ------------------------------
--------------------
/cluster/metrocluster       metrocluster configuration-settings connection
connect, metrocluster configuration-settings dr-group create, metrocluster
configuration-settings interface create, metrocluster configuration-
settings mediator add, metrocluster configure, metrocluster show, storage
aggregate create, storage aggregate mirror
                                                            GET, PATCH,
POST, DELETE
/cluster/metrocluster/diagnostics
                            metrocluster check show         GET, POST,
DELETE
/cluster/metrocluster/dr-groups
                            metrocluster configuration-settings dr-group
create, metrocluster configuration-settings dr-group show
                                                            GET, PATCH,
POST, DELETE
/cluster/metrocluster/interconnects
                            metrocluster interconnect mirror show
                                                            GET, PATCH
/cluster/metrocluster/nodes metrocluster node show          GET
/cluster/metrocluster/operations
                            metrocluster operation show     GET
/cluster/metrocluster/svms  metrocluster vserver show       GET
7 entries were displayed.
```

```
cluster1::> security login role show-rest -endpoint *cifs/session*
REST Endpoint               CLI Commands                   Supported
Operations
--------------------------- ------------------------------
--------------------
/protocols/cifs/session/files
                            vserver cifs session file      GET
/protocols/cifs/session/files/$id/$id
                            vserver cifs session file      GET, DELETE
/protocols/cifs/sessions    vserver cifs session           GET, POST,
DELETE
/protocols/cifs/sessions/$id/$id
                            vserver cifs session           GET, POST,
DELETE
4 entries were displayed.
```

security login role show

Show access control roles

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login role show` command displays the following information about access-control roles:

- Role name
- Command or command directory to which the role has access
- Access level (none, read-only, or all)
- Query (detailed view only)

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.

[-vserver <vserver name>] - Vserver

Selects the roles that match this parameter value.

[-role <text>] - Role Name

Selects the roles that match this parameter value. If this parameter and the `-cmddirname` parameter are both used, the command displays detailed information about the specified access-control role.

[-cmddirname <text>] - Command / Directory

Selects the roles that match this parameter value. If this parameter and the `-role` parameter are both used, the command displays detailed information about the specified access-control role.

[-access {none|readonly|read_create|read_modify|read_create_modify|all}] - Access Level

Selects the roles that match this parameter value.

[-query <query>] - Query

Selects the roles that match this parameter value.

Examples

The example below displays information about all access-control roles:

```
cluster1::> security login role show

Vserver    RoleName      Command/Directory                Query
AccessLevel
---------- ------------- -------------------------------- -----
-----------
vs         vsadmin       DEFAULT                                 none
vs         vsadmin       dashboard health vserver                readonly
vs         vsadmin       job                                     readonly
vs         vsadmin       job schedule                            none
vs         vsadmin       lun                                     all
vs         vsadmin       network connections                     readonly
cluster1   admin         DEFAULT                                 all
cluster1   readonly      DEFAULT                                 readonly
cluster1   readonly      volume                                  none
```

security login role config modify

Modify local user account restrictions

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login role config modify` command modifies user account and password restrictions.

For the password character restrictions documented below (uppercase, lowercase, digits, etc.), the term "characters" refers to ASCII-range characters only - not extended characters.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver name associated with the profile configuration.

`-role <text>` - Role Name

This specifies the role whose account restrictions are to be modified.

`[-username-minlength <integer>]` - Minimum Username Length Required

This specifies the required minimum length of the user name. Supported values are 3 to 16 characters. The default setting is *3* characters.

`[-username-alphanum {enabled|disabled}]` - Username Alpha-Numeric

This specifies whether a mix of alphabetic and numeric characters are required in the user name. If this parameter is enabled, a user name must contain at least one letter and one number. The default setting is *disabled*.

`[-passwd-minlength <integer>]` - Minimum Password Length Required

This specifies the required minimum length of a password. Supported values are 3 to 64 characters. The default setting is *8* characters.

`[-passwd-alphanum {enabled|disabled}]` - Password Alpha-Numeric

This specifies whether a mix of alphabetic and numeric characters is required in the password. If this parameter is enabled, a password must contain at least one letter and one number. The default setting is *enabled*.

`[-passwd-min-special-chars <integer>]` - Minimum Number of Special Characters Required in the Password

This specifies the minimum number of special characters required in a password. Supported values are from 0 to 64 special characters. The default setting is *0* , which requires no special characters.

`[-passwd-expiry-time <integer_or_unlimited>]` - Password Expires In (Days)

This specifies password expiration in days. A value of *0* means all passwords associated with the accounts in the role expire now. The default setting is *unlimited* , which means the passwords never expire.

`[-require-initial-passwd-update {enabled|disabled}]` - Require Initial Password Update on First Login

This specifies whether users must change their passwords when logging in for the first time. Initial password changes can be done only through SSH or serial-console connections. The default setting is *disabled*.

`[-max-failed-login-attempts <integer>]` - Maximum Number of Failed Attempts

This specifies the allowed maximum number of consecutive invalid login attempts. When the failed login attempts reach the specified maximum, the account is automatically locked. The default is *5* , which means *5* failed login attempts will cause an account to be locked.

`[-lockout-duration <integer>]` - (DEPRECATED)-Maximum Lockout Period (Days)

(DEPRECATED)-This specifies the number of days for which an account is locked if the failed login attempts reach the allowed maximum. The default is *0* , which means the accounts will be locked for *1* hour. For roles which were created in a release before ONTAP 9.15.0 with the default value of *0* , this value will be automatically changed to *1* during upgrade to ONTAP 9.15.0. In other words, the value of this field for roles created before ONTAP 9.15.0 is defaulted to 24 hrs. For the roles which are created in ONTAP 9.15.0 or later, the value of this field defaults to 1 hour. This parameter is deprecated in ONTAP 9.15.0 and later. It may be removed from a future release of ONTAP.

[-disallowed-reuse <integer>] - Disallow Last 'N' Passwords

This specifies the number of previous passwords that are disallowed for reuse. The default setting is six, meaning that the user cannot reuse any of their last six passwords. The minimum allowed value is *6* .

[-change-delay <integer>] - Delay Between Password Changes (Days)

This specifies the number of days that must pass between password changes. The default setting is *0* .

[-delay-after-failed-login <integer>] - Delay after Each Failed Login Attempt (Secs)

This specifies the amount of delay observed by the system in seconds upon invalid login attempts. The default setting is *4* seconds.

[-passwd-min-lowercase-chars <integer>] - Minimum Number of Lowercase Alphabetic Characters Required in the Password

This specifies the minimum number of lowercase characters required in a password. Supported values are from 0 to 64 lowercase characters. The default setting is *0* , which requires no lowercase characters.

[-passwd-min-uppercase-chars <integer>] - Minimum Number of Uppercase Alphabetic Characters Required in the Password

This specifies the minimum number of uppercase characters required in a password. Supported values are from 0 to 64 uppercase characters. The default setting is *0* , which requires no uppercase characters.

[-passwd-min-digits <integer>] - Minimum Number of Digits Required in the Password

This specifies the minimum number of digits required in a password. Supported values are from 0 to 64 digits characters. The default setting is *0* , which requires no digits.

[-passwd-expiry-warn-time <integer_or_unlimited>] - Display Warning Message Days Prior to Password Expiry (Days)

This specifies the warning period for password expiry in days. A value of *0* means warn user about password expiry upon every successful login. The default setting is *unlimited* , which means never warn about password expiry.

[-account-expiry-time <integer_or_unlimited>] - Account Expires in (Days)

This specifies account expiration in days. The default setting is *unlimited* , which means the accounts never expire. The account expiry time must be greater than account inactive limit.

[-account-inactive-limit <integer_or_unlimited>] - Maximum Duration of Inactivity before Account Expiration (Days)

This specifies inactive account expiry limit in days. The default setting is *unlimited* , which means the inactive accounts never expire. The account inactive limit must be less than account expiry time.

[-account-lockout-duration {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W | disabled}] - Account Lockout Duration (ISO 8601 Duration Format)

This specifies the duration in ISO 8601 format for which an account is locked if the failed login attempts reach the allowed maximum. The default is 1 hour. Specifying this field will set the field `lockout-duration` as *0* and the field `account-lockout-duration` will be used for further operations.

Examples

The following command modifies the user-account restrictions for an account with the role name admin for a Vserver named vs1. The minimum size of the password is set to 12 characters.

```
cluster1::> security login role config modify -role admin -vserver vs1
    -passwd-minlength 12
```

The following command sets the maximum allowed number of consecutive invalid login attempts to 3 and the maximum account lockout duration to 1 minute 30 seconds after 3 failed login attempts for role admin for vserver vs1:

```
cluster1::> security login role config modify -role admin -vserver vs1
    -max-failed-login-attempts 3 -account-lockout-duration PT1M30S
```

security login role config reset

Reset RBAC characteristics supported on releases later than Data ONTAP 8.1.2

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security login role config reset` command resets the following role based access control (RBAC) characteristics to their default values. The system prompts you to run this command if you revert to ONTAP 8.1.2 or earlier. If you do not reset these characteristics, the revert process will fail.

- Minimum number of special characters required in password ("0")
- Minimum number of uppercase characters required in password ("0")
- Minimum number of lowercase characters required in password ("0")
- Minimum number of digits required in password ("0")
- Password-expiration time, in days ("unlimited")
- Whether the password must be changed at the initial login ("disabled")
- Maximum number of failed login attempts permitted before the account is locked out ("5")
- Maximum time period for which the user account is locked out after the maximum number of failed login attempts is reached ("1 Hour")

Examples

The following command resets the above mentioned RBAC characteristics of all cluster and Vserver roles to their default values.

```
cluster1::> security login role config reset
```

security login role config show

Show local user account restrictions

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security login role config show` command displays the following information about account restrictions for management-utility user accounts:

- Role name `-role`
- Minimum size of the password, in characters `-passwd-minlength`
- Whether the password requires alphanumeric characters `-passwd-alphanum`
- Number of previous passwords that cannot be reused `-disallowed-reuse`
- Minimum number of days that must elapse before users can change their passwords `-change-delay`

You can display detailed information about the restrictions on a specific account by specifying the `-role` parameter. This adds the following information:

- Minimum length of the user name, in characters `-username-minlength`
- Whether the user name requires alphanumeric characters `-username-alphanum`
- Minimum length of the password, in characters `-passwd-minlength`
- Whether the password requires alphanumeric characters `-passwd-alphanum`
- Minimum number of special characters required in password `-passwd-min-special-chars`
- Minimum number of lowercase characters required in password `-passwd-min-lowercase-chars`
- Minimum number of uppercase characters required in password `-passwd-min-uppercase-chars`
- Minimum number of digits required in password `-passwd-min-digits`
- Minimum number of days that must elapse before users can change their passwords `-change-delay`
- Whether the password must be changed at the initial login `-require-initial-passwd-update`
- Password-expiration time, in days `-passwd-expiry-time`
- Display warning message days prior to password expiry `-passwd-expiry-warn-time`
- Number of previous passwords that cannot be reused `-disallowed-reuse`
- Maximum number of failed login attempts permitted before the account is locked out `-max-failed -login-attempts`
- (DEPRECATED)-Number of days for which the user account is locked after the maximum number of failed login attempts is reached. For roles which were created in a release before ONTAP 9.15.0 with the default value of *`0`* , this value will be automatically changed to *`1`* during upgrade to ONTAP 9.15.0. In other words, the value of this field for roles created before ONTAP 9.15.0 is defaulted to 24 hrs. For the roles which are created in ONTAP 9.15.0 or later, the value of this field defaults to 1 hour. This parameter is deprecated in ONTAP 9.15.0 and later. It may be removed from a future release of ONTAP `-lockout-duration`
- Account-expiration time, in days `-account-expiry-time`
- Maximum duration of inactivity before account expiration, in days `-account-inactive-limit`
- Delay after each failed login attempt, in secs `-delay-after-failed-login`
- Duration for which the user account is locked after the maximum number of failed login attempts is reached `-account-lockout-duration`

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.

[-vserver <vserver name>] - Vserver

Selects the profile configurations that match this parameter value

[-role <text>] - Role Name

If this parameter is specified, the command displays detailed information about restrictions for the specified user account.

[-username-minlength <integer>] - Minimum Username Length Required

Selects the profile configurations that match this parameter value.

[-username-alphanum {enabled|disabled}] - Username Alpha-Numeric

Selects the profile configurations that match this parameter value. Enabled means a user name must contain both letters and numbers.

[-passwd-minlength <integer>] - Minimum Password Length Required

Selects the profile configurations that match this parameter value.

[-passwd-alphanum {enabled|disabled}] - Password Alpha-Numeric

Selects the profile configurations that match this parameter value. Enabled means a password must contain both letters and numbers.

[-passwd-min-special-chars <integer>] - Minimum Number of Special Characters Required in the Password

Selects the profile configurations that match this parameter value.

[-passwd-expiry-time <integer_or_unlimited>] - Password Expires In (Days)

Selects the profile configurations that match this parameter value.

[-require-initial-passwd-update {enabled|disabled}] - Require Initial Password Update on First Login

Selects the profile configurations that match this parameter value.

[-max-failed-login-attempts <integer>] - Maximum Number of Failed Attempts

Selects the profile configurations that match this parameter value.

[-lockout-duration <integer>] - (DEPRECATED)-Maximum Lockout Period (Days)

Selects the profile configurations that match this parameter value.

[-disallowed-reuse <integer>] - Disallow Last 'N' Passwords

Selects the profile configurations that match this parameter value.

[-change-delay <integer>] - Delay Between Password Changes (Days)

Selects the profile configurations that match this parameter value.

[-delay-after-failed-login <integer>] - Delay after Each Failed Login Attempt (Secs)

Selects the profile configurations that match this parameter value.

[-passwd-min-lowercase-chars <integer>] - Minimum Number of Lowercase Alphabetic Characters Required in the Password

Selects the profile configurations that match this parameter value.

[-passwd-min-uppercase-chars <integer>] - Minimum Number of Uppercase Alphabetic Characters Required in the Password

Selects the profile configurations that match this parameter value.

[-passwd-min-digits <integer>] - Minimum Number of Digits Required in the Password

Selects the profile configurations that match this parameter value.

[-passwd-expiry-warn-time <integer_or_unlimited>] - Display Warning Message Days Prior to Password Expiry (Days)

Selects the profile configurations that match this parameter value.

[-account-expiry-time <integer_or_unlimited>] - Account Expires in (Days)

Selects the profile configurations that match this parameter value.

[-account-inactive-limit <integer_or_unlimited>] - Maximum Duration of Inactivity before Account Expiration (Days)

Selects the profile configurations that match this parameter value.

[-account-lockout-duration {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W | disabled}] - Account Lockout Duration (ISO 8601 Duration Format)

Selects the profile configurations that match this parameter value.

Examples

The example below displays restriction information about all user accounts:

```
cluster1::> security login role config show
                           ----- Password Restrictions -----
Vserver     RoleName      Size AlphaNum NoReuse ChangeDelay
----------- ------------- ---- -------- ------- -----------
vs          vsadmin          8  enabled       6      0 days
vs          vsadmin-protocol 8  enabled       6      0 days
vs          vsadmin-readonly 8  enabled       6      0 days
vs          vsadmin-volume   8  enabled       6      0 days
cluster1    admin            6  enabled       6      0 days
cluster1    readonly         6  enabled       6      0 days
```

security login totp create

Add a TOTP secret

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login totp create` associates an new secret key with a user account. This command requires that you enter a user name.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver of the user for whom you are adding the secret key.

`-username <text>` - Username

This parameter specifies the name of the user for whom you are adding the secret key. If you do not specify a user in case of Cserver, the user named admin is specified by default. If it is a data vserver, it defaults to logged-in user

`[-comment <text>]` - Comment

This optionally specifies comment text for the TOTP key. Note that comment text should be enclosed in quotation marks.

Examples

The following command creates a secret key with a user named tsmith for Vserver vs1. The secret key has a comment text is “This is a new key”.

```
cluster1::> security login totp create -vserver vs1 -username tsmith
-comment "This is a new key"
```

security login totp delete

Delete a TOTP secret

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login totp delete` command deletes secret key for a specific user. To delete a secret key, you must specify a user name.

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver of the user for whom you are deleting the secret key.

`-username <text>` - Username

This parameter specifies the name of the user for whom you are deleting a secret key. If you do not specify a user in case of cserver, the user named admin is specified by default. In case of data vserver, it defaults to logged-in user

Examples

The following command deletes the secret key for the user named tsmith.

```
cluster1::> security login totp delete -username tsmith -index 5
```

security login totp modify

Modify a TOTP status

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login totp modify` command modifies the TOTP status. Only admin can modify the TOTP status and it requires that you enter a user name and enabled

Parameters

`-vserver <vserver name>` - Vserver

This parameter optionally specifies the Vserver of the user for whom you are modifying the TOTP status or comment.

`-username <text>` - Username

This parameter specifies the name of the user for whom you are modifying the TOTP status.

`[-enabled {true|false}]` - TOTP Status

This parameter specifies the TOTP status of the user for whom you are modifying.

`[-comment <text>]` - Comment

This optionally specifies comment text for the TOTP key. Note that comment text should be enclosed in quotation marks.

Examples

The following command disable the TOTP for the user named tsmith for Vserver vs1.

```
cluster1::> security login totp modify -vserver vs1 -username tsmith
-enabled false
```

security login totp show

Display TOTP secret

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security login totp show` command displays information about secret keys.

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.

[-vserver <vserver name>] - Vserver

Selects the secret keys that match this parameter value.

[-username <text>] - Username

Selects the secret keys that match this parameter value.

[-fingerprint <text>] - TOTP Secret key fingerprint

Select the secret keys that match this parameter value.

[-enabled {true|false}] - TOTP Status

Selects the secret keys that match this parameter value.

[-comment <text>] - Comment

Selects the secret keys that match this parameter value.

Examples

The example below displays secret key information for the user named tsmith.

```
cluster1::> security login totp show -username tsmith
UserName: tsmith
TOTP SHA Fingerprint:
7038eb494f8b86726bdfae9da1fbadb348f8fe26116f49e09f718a7b8bdd73b8
TOTP status: true
Comment: This is a new key
```

security multi-admin-verify commands

security multi-admin-verify modify

Modify multi-admin-verify settings

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify modify` command is used to modify the Multi-Admin-Verify global settings.

Parameters

`[-approval-groups <text>,…]` - List of Global Approval Groups

This specifies the list of global approval groups which are inherited by the rule if the approval-groups is not provided for the rule. The default value is an empty list. The approval-groups should be defined to enable multi-admin verification. The supplied value replaces the list. You can create an approval-group by using the security multi-admin-verify approval-group create command.

`[-required-approvers <integer>]` - Number of Required Approvers

This specifies the required number of approvers to approve the request which is inherited by the rule if required-approvers is not provided for the rule. The default and minimum number of required approvers is 1.

`[-enabled {true|false}]` - Is Multi-Admin-Verify Enabled

This specifies the current state. Multi-admin verification is not required to enable the feature. However, it is required to disable the feature. By default, the feature is disabled and the value is set to false. It is recommended that multi-admin-verify is enabled equally on peered ONTAP clusters.

`[-execution-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>]` - Execution Expiry

This is the amount of time that the authorized users have after a request is approved to execute the requested operation before the request expires. The default value is one hour (*`1h`*), the minimum supported value is one second (*`1s`*), and the maximum supported value is 14 days (*`14d`*).

`[-approval-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>]` - Approval Expiry

This is the amount of time that the approvers have after a new execution request is submitted to approve or disapprove the request before the request expires. The default value is one hour (*`1h`*), the minimum supported value is one second (*`1s`*), and the maximum supported value is 14 days (*`14d`*).

Examples

This command changes the approval groups:

```
cluster1::> security multi-admin-verify modify -approval-groups group1,
group2
```

This command changes the required number of approvers:

```
cluster1::> security multi-admin-verify modify -required-approvers 3
```

This command enables the feature. The default is false (disabled):

```
cluster1::> security multi-admin-verify modify -enabled true
```

This command changes the execution expiry:

```
cluster1::> security multi-admin-verify modify -execution-expiry 14d
```

This command changes the approval expiry:

```
cluster1::> security multi-admin-verify modify -approval-expiry 48h
```

Related Links

- security multi-admin-verify approval-group create

security multi-admin-verify show

Display multi-admin-verify configuration

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify show` command displays the object store that contains the global setting values of the multi-admin-verify feature.

- Is Enabled: Displays the current state of the feature. This feature is, by default, disabled and the value is set to false.
- Required Approvers: Displays the required number of approvers to approve the ONTAP execution request. This is inherited by the rule if required-approvers is not provided for the rule. The default and minimum number of required approvers is 1.
- Approval Expiry: Displays the amount of time that the approvers have after a new execution request is submitted to approve or disapprove the request before the request expires.
- Execution Expiry: Displays the amount of time that the authorized users have after a request is approved to execute the requested operation before the request expires.
- Approval Groups: Displays the list of global approval groups. This will be in effect if the approval groups is not specified for a multi-admin-verify rule.

Examples

The following example displays typical global settings information:

```
cluster1::> security multi-admin-verify show
Is      Required  Execution Approval Approval
    Enabled Approvers Expiry    Expiry   Groups
    ------- --------- --------- -------- -----------
    false   1         1h        1h       group1, group2
```

security multi-admin-verify approval-group create

Create an Approval Group

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify approval-group create` command creates an approval group for a specified Vserver for a specified list of ONTAP users.

Parameters

`-vserver <vserver>` - Vserver

This specifies the Vserver name to which the approval group is associated with. This is an optional parameter. This parameter defaults to a Cluster server and supports only Cluster servers.

`-name <text>` - Group Name

This specifies the name of an approval group.

`-approvers <text>,…` - Approvers

This specifies the list of ONTAP users that are part of the approval group. Each specified user must belong to the specified Vserver.

`[-email <mail address>,…]` - Email Addresses

This specifies the email addresses that are notified when a request is created, approved, vetoed, or executed.

Examples

The following example creates a new approval group named group1 with approver admin1 that is associated with the default Vserver cluster1:

```
cluster1::> security multi-admin-verify approval-group create -name group1
-approvers admin1
```

security multi-admin-verify approval-group delete

Delete an Approval Group

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify approval-group delete` command deletes the specified approval group.

Parameters

`-vserver <vserver>` - Vserver

This specifies the Vserver information to which the approval group is associated with. This is an optional parameter. This parameter defaults to Cluster server and supports only Cluster servers.

`-name <text>` - Group Name

This specifies the name of an approval group to be deleted.

Examples

The following example deletes the approval group, group1:

```
cluster1::> security multi-admin-verify approval-group delete -name group1
```

security multi-admin-verify approval-group modify

Modify an Approval Group

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify approval-group modify` command is used to modify attributes of an approval group.

Parameters

`-vserver <vserver>` - Vserver

This specifies the Vserver name to which the approval group is associated with. This is an optional parameter. This parameter defaults to a Cluster server and supports only Cluster servers.

`-name <text>` - Group Name

This specifies the name of an approval group.

`[-approvers <text>,…]` - Approvers

This specifies the list of ONTAP users that are part of the approval group. Each specified user must belong to the specified Vserver.

`[-email <mail address>,…]` - Email Addresses

This specifies the email addresses that are notified when a request is created, approved, vetoed, or executed.

Examples

This command changes the approvers:

```
cluster1::> security multi-admin-verify approval-group modify -name group1
-approvers admin1
```

security multi-admin-verify approval-group replace

Add and/or remove approvers from the list

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify approval-group replace` command is used to replace the list of approvers of an approval group.

Parameters

`-vserver <vserver>` - Vserver

This specifies the Vserver name to which the approval group is associated with. This is an optional parameter. This parameter defaults to Cluster server and supports only Cluster servers.

`-name <text>` - Group Name

This specifies the name of the approval group whose approvers are to be replaced.

`[-approvers-to-add <text>,…]` - New Approvers

This specifies the list of ONTAP users that are to be added to the current list of approvers of the approval group.

`[-approvers-to-remove <text>,…]` - Existing Approvers

This specifies the list of ONTAP users that are to be removed from the current list of approvers of the approval group.

Examples

The following example adds user admin2 and removes user admin from the current approvers list, while group1 is associated with the default Vserver:

```
cluster1::> security multi-admin-verify approval-group replace -name
group1 -approvers-to-add admin2 -approvers-to-remove admin.
```

security multi-admin-verify approval-group show

Display Approval Groups

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify approval-group show` command displays information about approval groups and the users that are registered with each group.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You can use '-fields ?' to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver>] - Vserver

This specifies the Vserver name to which the approval group is associated with. This is an optional parameter. This parameter defaults to a Cluster server and supports only Cluster servers.

[-name <text>] - Group Name

This specifies the name of an approval group.

[-approvers <text>,…] - Approvers

This specifies the list of ONTAP users that are part of the approval group. Each specified user must belong to the specified Vserver.

[-email <mail address>,…] - Email Addresses

This specifies the email addresses that are notified when a request is created, approved, vetoed, or executed.

Examples

The following example displays typical approval groups information:

```
cluster1::> security multi-admin-verify approval-group show
    Vserver  Name             Approvers
    -------- ----------------
-------------------------------------------------------
    cluster1
             group1           admin
             group2           admin, admin1
    2 entries were displayed.
```

security multi-admin-verify request approve

Approve a request

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify request approve` command approves the specified request. Cluster peering commands might require remote approvals.

Parameters

`-index <integer>` - Request Index

This allows the user to specify the index of the request to be approved.

`[-execute-on-approval {true|false}]` - Execute Command On Final Approval

If set to *`true`* , the operation will automatically execute on final approval.

Examples

The following example approves the request with index 1:

```
cluster1::> security multi-admin-verify request approve -index 1
```

security multi-admin-verify request create

Create a request

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify request create` command creates a request for the specified ONTAP operation.

Parameters

`[-index <integer>]` - Request Index

This specifies the index of the request which is automatically generated for each request.

`-operation <text>` - Operation

This specifies the ONTAP operation information for which the request is to be created.

`-query <query>` - Query

This identifies the object (or objects) upon which the user wants to apply the operation. Any field or query supported by the operation can be supplied.

`[-comment <text>]` - Comment

This is an optional parameter where users creating a request can provide comments related to the request.

`[-users-permitted <text>,…]` - Users Permitted

This is an optional parameter where a user creating the request can specify the list of ONTAP users who are permitted to perform the ONTAP operation specified by the request, once it is approved. If this parameter is not provided, then any user with default permissions to perform the ONTAP operation is allowed to perform the ONTAP operation specified by the request.

Examples

The following example creates a new request for ONTAP operation volume delete which is applicable to objects of vserver vs0.

```
cluster1::> security multi-admin-verify request create -operation "volume
delete" -query "-vserver vs0"
```

The following example creates a new request for the ONTAP operation volume snapshot delete which is applicable to Vserver objects vs0 and volume v1. Users permitted to perform this operation on the specified subset of objects are user1 and user2:

```
cluster1::> security multi-admin-verify request create -operation "volume
delete" -query "-vserver vs0 -volume v1" -users-permitted user1, user2
```

security multi-admin-verify request delete

Delete a request

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify request delete` command deletes the specified request.

Parameters

`-index <integer>` - Request Index

This allows the user to specify the index of the request to be deleted.

Examples

The following example deletes the request with index 1:

```
cluster1::> security multi-admin-verify request delete -index 1
```

security multi-admin-verify request show-pending

Show only pending requests

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify request show-pending` command displays information about multi-admin verification requests that are in the pending state.

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.

[-index <integer>] - Request Index

This specifies the index of the request.

[-operation <text>] - Operation

This specifies the ONTAP operation for which the request is created.

[-query <query>] - Query

This identifies the object (or objects) upon which the user wants to apply the operation.

[-required-approvers <integer>] - Required Approvers

This specifies the number of distinct users that are required to approve the request. A user can set the required-approvers to the ONTAP operation rule. If a user does not set the required-approvers to the rule, then the required-approvers from the global setting is applied.

[-pending-approvers <integer>] - Pending Approvers

This specifies the number of distinct users that are still required to approve the request for the request to be marked as approved.

[-approval-expiry {MM/DD/YYYY HH:MM:SS | {1..7}days | {1..168}hours | PnDTnHnMnS | PnW}] - Approval Expiry

This specifies the expiry information within which an approve or veto action is to be taken by the approvers from the time the request is submitted. Any authorized user can set the approval-expiry to the ONTAP operation rule. If the user does not set the approval-expiry to the rule, then the approval-expiry from the global setting is applied.

[-execution-expiry {MM/DD/YYYY HH:MM:SS | {1..7}days | {1..168}hours | PnDTnHnMnS | PnW}] - Execution Expiry

This specifies the expiry information within which an ONTAP operation is to be executed from the time the request is approved. An authorized user can set the execution-expiry to the ONTAP operation rule. If the user does not set the execution-expiry to the rule, then the execution-expiry from the global setting is applied.

[-users-approved <text>,…] - Approvals

This specifies the list of users that have approved the request.

[-user-vetoed <text>] - User Vetoed

This specifies the user who vetoed the request.

[-vserver <vserver>] - Vserver

This specifies the Vserver information to which the request is associated with.

[-user-requested <text>] - User Requested

This specifies the username who created the request.

[-time-created <MM/DD/YYYY HH:MM:SS>] - Time Created

This specifies the time at which the request is created.

[-time-approved <MM/DD/YYYY HH:MM:SS>] - Time Approved

This specifies the time at which the request state changed to approved.

[-comment <text>] - Comment

This specifies the comments that are associated with the request.

[-users-permitted <text>,…] - Users Permitted

This specifies the list of users that are permitted to perform the ONTAP operation for which the request is approved. If users-permitted is empty, then any user who, by default, has permission to perform the ONTAP operation is allowed.

[-execute-on-approval {true|false}] - Execute Approved Command

This specifies whether the operation being approved is automatically executed. If set, the operation is executed immediately when the request is is marked as approved.

Examples

The following example displays typical request information:

```
cluster1::> security multi-admin-verify request show-pending
Pending
    Index Operation              Query                   State
Approvers Requestor
    ----- ---------------------- ----------------------- --------
--------- ---------
        1 volume delete                                  pending  3
admin
```

security multi-admin-verify request show

Display requests

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify request show` command displays information about multi-admin verification requests.

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.

[-index <integer>] - Request Index

This specifies the index of the request.

[-operation <text>] - Operation

This specifies the ONTAP operation for which the request is created.

[-query <query>] - Query

This identifies the object (or objects) upon which the user wants to apply the operation.

[-state {pending|approved|vetoed|expired|executed}] - State

This specifies the query information that is applied to the subset of objects of ONTAP operation of the request.

[-required-approvers <integer>] - Required Approvers

This specifies the number of distinct users that are required to approve the request. A user can set the required-approvers to the ONTAP operation rule. If a user does not set the required-approvers to the rule, then the required-approvers from the global setting is applied.

[-pending-approvers <integer>] - Pending Approvers

This specifies the number of distinct users that are still required to approve the request for the request to be marked as approved.

[-approval-expiry {MM/DD/YYYY HH:MM:SS | {1..7}days | {1..168}hours | PnDTnHnMnS | PnW}] - Approval Expiry

This specifies the expiry information within which an approve or veto action is to be taken by the approvers from the time the request is submitted. Any authorized user can set the approval-expiry to the ONTAP operation rule. If the user does not set the approval-expiry to the rule, then the approval-expiry from the global setting is applied.

[-execution-expiry {MM/DD/YYYY HH:MM:SS | {1..7}days | {1..168}hours | PnDTnHnMnS | PnW}] - Execution Expiry

This specifies the expiry information within which an ONTAP operation is to be executed from the time the request is approved. An authorized user can set the execution-expiry to the ONTAP operation rule. If the user does not set the execution-expiry to the rule, then the execution-expiry from the global setting is applied.

[-users-approved <text>,…] - Approvals

This specifies the list of users that have approved the request.

[-user-vetoed <text>] - User Vetoed

This specifies the user who vetoed the request.

[-vserver <vserver>] - Vserver

This specifies the Vserver information to which the request is associated with.

[-user-requested <text>] - User Requested

This specifies the username who created the request.

[-time-created <MM/DD/YYYY HH:MM:SS>] - Time Created

This specifies the time at which the request is created.

[-time-approved <MM/DD/YYYY HH:MM:SS>] - Time Approved

This specifies the time at which the request state changed to approved.

[-comment <text>] - Comment

This specifies the comments that are associated with the request.

[-users-permitted <text>,…] - Users Permitted

This specifies the list of users that are permitted to perform the ONTAP operation for which the request is approved. If users-permitted is empty, then any user who, by default, has permission to perform the ONTAP operation is allowed.

[-execute-on-approval {true|false}] - Execute Approved Command

This specifies whether the operation being approved is automatically executed. If set, the operation is executed immediately when the request is is marked as approved.

Examples

The following example displays typical request information:

```
cluster1::> security multi-admin-verify request show
Pending
    Index Operation              Query                   State
Approvers Requestor
    ----- ---------------------- ------------------------ --------
--------- ---------
        1 volume delete                                  pending  3
admin
```

security multi-admin-verify request veto

Veto a request

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify request veto` command vetoes the specified request.

Parameters

`-index <integer>` - Request Index

This allows the user to specify the index of the request to be vetoed.

Examples

The following example vetoes the request with index 1:

```
cluster1::> security multi-admin-verify request veto -index 1
```

security multi-admin-verify rule create

Create a rule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify rule create` command creates a rule for the specified ONTAP operation.

Parameters

`[-vserver <vserver>]` - Vserver

This specifies Vserver information for which the rule should be associated with. This is an optional parameter. This parameter defaults to a Cluster server and supports only Cluster servers.

`-operation <text>` - Operation

This specifies the ONTAP operation information for the rule to be created.

`[-auto-request-create {true|false}]` - Automatic Request Creation

This specifies rule information for the auto request create state. Auto request creation for the rule is enabled by default, by setting this value to true.

`[-query <query>]` - Query

This specifies the query information which is applied to the subset of objects of ONTAP operation of the rule to be created. This is an optional parameter. If a query is not specified for the rule, the rule applies to all objects of the ONTAP operation.

`[-required-approvers {<integer>|-}]` - Required Number of Approvers

This specifies the required number of approvers to approve the ONTAP execution request. This is an optional parameter. If required-approvers is not specified for the rule, the required-approvers from the global setting is applied to the ONTAP operation request. The required-approvers from the global setting can be viewed using the security multi-admin-verify show command. The minimum supported value is 1.

`[-approval-groups <text>,…]` - Approval Groups

This specifies the list of users who can approve the ONTAP operation request. This is an optional parameter. If approval-groups is not specified for the rule, the approval-groups from the global setting is applied to the ONTAP operation request. The approval-groups from the global setting can be viewed using

the security multi-admin-verify show command.

[-execution-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Execution Expiry

This specifies the amount of time after a request has been approved by which the operation must be executed before the approved execution request expires. This is an optional parameter. If execution-expiry is not specified for the rule, the execution-expiry from the global setting is applied to the ONTAP execution request. The execution-expiry from the global setting can be viewed using the security multi-admin-verify show command. The default value is one hour (*1h*), the minimum supported value is one second (*1s*), and the maximum supported value is 14 days (*14d*).

[-approval-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Approval Expiry

This specifies the amount of time after a new execution request is submitted by which approvers have to approve or disapprove the request before the pending execution request expires. This is an optional parameter. If approval-expiry is not specified for the rule, the approval-expiry from the global setting is applied to the ONTAP execution request. The approval-expiry from the global setting can be viewed using the security multi-admin-verify show command. The default value is one hour (*1h*), the minimum supported value is one second (*1s*), and the maximum supported value is 14 days (*14d*).

Examples

The following example creates a new rule for the ONTAP operation volume delete with 3 required approvers and is applicable to Vserver vs0 objects:

```
cluster1::> security multi-admin-verify rule create  -operation "volume
delete" -query "-vserver vs0" -required-approvers 3
```

Related Links

- security multi-admin-verify show

security multi-admin-verify rule delete

Delete a rule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify rule delete` command deletes the specified rule.

Parameters

-vserver <vserver> - Vserver

This specifies the Vserver information to which the rule is associated with. This is an optional parameter. This parameter defaults to a Cluster server and supports only Cluster servers.

-operation <text> - Operation

This specifies the ONTAP operation whose associated rule is to be deleted.

Examples

The following example deletes the rule for ONTAP operation volume delete and the default Vserver cluster1:

```
cluster1::> security multi-admin-verify rule delete -operation "volume
delete"
```

security multi-admin-verify rule modify

Modify a rule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify rule modify` command is used to modify the attributes of the rule.

Parameters

`-vserver <vserver>` - Vserver

This specifies Vserver information for which the rule should be associated with. This is an optional parameter. This parameter defaults to a Cluster server and supports only Cluster servers.

`-operation <text>` - Operation

This specifies the ONTAP operation information for the rule to be created.

`[-auto-request-create {true|false}]` - Automatic Request Creation

This specifies rule information for the auto request create state. Auto request creation for the rule is enabled by default, by setting this value to true.

`[-query <query>]` - Query

This specifies the query information which is applied to the subset of objects of ONTAP operation of the rule to be created. This is an optional parameter. If a query is not specified for the rule, the rule applies to all objects of the ONTAP operation.

`[-required-approvers {<integer>|-}]` - Required Number of Approvers

This specifies the required number of approvers to approve the ONTAP execution request. This is an optional parameter. If required-approvers is not specified for the rule, the required-approvers from the global setting is applied to the ONTAP operation request. The required-approvers from the global setting can be viewed using the security multi-admin-verify show command. The minimum supported value is 1.

`[-approval-groups <text>,…]` - Approval Groups

This specifies the list of users who can approve the ONTAP operation request. This is an optional parameter. If approval-groups is not specified for the rule, the approval-groups from the global setting is applied to the ONTAP operation request. The approval-groups from the global setting can be viewed using the security multi-admin-verify show command.

`[-execution-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>]` - Execution Expiry

This specifies the amount of time after a request has been approved by which the operation must be

executed before the approved execution request expires. This is an optional parameter. If execution-expiry is not specified for the rule, the execution-expiry from the global setting is applied to the ONTAP execution request. The execution-expiry from the global setting can be viewed using the security multi-admin-verify show command. The default value is one hour (*1h*), the minimum supported value is one second (*1s*), and the maximum supported value is 14 days (*14d*).

[-approval-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Approval Expiry

This specifies the amount of time after a new execution request is submitted by which approvers have to approve or disapprove the request before the pending execution request expires. This is an optional parameter. If approval-expiry is not specified for the rule, the approval-expiry from the global setting is applied to the ONTAP execution request. The approval-expiry from the global setting can be viewed using the security multi-admin-verify show command. The default value is one hour (*1h*), the minimum supported value is one second (*1s*), and the maximum supported value is 14 days (*14d*).

Examples

This command changes the approval groups:

```
cluster1::> security multi-admin-verify rule modify -operation "volume
delete" -approval-groups group1, group2
```

This command changes the required number of approvers:

```
cluster1::> security multi-admin-verify rule modify -operation "volume
snapshot delete" -required-approvers 3
```

This command changes the query:

```
cluster1::> security multi-admin-verify rule modify -operation "volume
delete" -query "-vserver vs1"
```

This command changes the execution expiry:

```
cluster1::> security multi-admin-verify rule modify -operation "volume
delete" -execution-expiry 14d
```

This command changes the approval expiry:

```
cluster1::> security multi-admin-verify rule modify -operation "volume
delete" -approval-expiry 48h
```

Related Links

- security multi-admin-verify show

security multi-admin-verify rule show

Display rules

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security multi-admin-verify rule show` command displays information about multi admin verification rules.

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.

[-vserver <vserver>] - Vserver

This specifies the Vserver information to which the rule is associated with. This is an optional parameter. This parameter defaults to a Cluster server and supports only Cluster servers.

[-operation <text>] - Operation

This specifies the ONTAP operation information for which the rule is created.

[-auto-request-create {true|false}] - Automatic Request Creation

This specifies the information of the auto request create state for the rule.

[-query <query>] - Query

This specifies the query information which is applied to the subset of objects of ONTAP operation of the rule.

[-required-approvers {<integer>|-}] - Required Number of Approvers

This specifies the number of approvers that are required to approve the ONTAP execution request.

[-approval-groups <text>,…] - Approval Groups

This specifies the list of approval groups that lists the users who can approve the ONTAP execution request.

[-execution-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Execution Expiry

This specifies the amount of time that the authorized users have after a request is approved to execute the requested operation before the request expires.

[-approval-expiry <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Approval Expiry

This is the amount of time that the approvers have after a new execution request is submitted to approve or disapprove the request before the request expires.

[-time-created <MM/DD/YYYY HH:MM:SS>] - Time Created

This specifies the time at which the rule is created.

[-system-defined {true|false}] - Is System Defined

Displays the value true if rule is defined by the system. Displays the value false if rule is defined by the user.

Examples

The following example displays typical rule information:

```
cluster1::> security multi-admin-verify rule show
                                                          Required
Approval
    Vserver     Operation                                 Approvers
Groups
    ----------- ----------------------------------------- ---------
-------------
    cluster1
                security login password                   1         -
                  Query: -multi-admin-approver true -different-user true
                security multi-admin-verify approval-group create
                                                          1         -
                security multi-admin-verify approval-group delete
                                                          1         -
                security multi-admin-verify approval-group modify
                                                          1         -
                security multi-admin-verify approval-group replace
                                                          1         -
                security multi-admin-verify modify        1         -
                security multi-admin-verify rule create   1         -
                security multi-admin-verify rule delete   1         -
                security multi-admin-verify rule modify   1         -
                volume delete                             3         -
                  Query: -vserver vs0
    10 entries were displayed.
```

security oauth2 commands

security oauth2 modify

Modify global OAuth 2.0 configuration

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security oauth2 modify` command enables or disables the OAuth 2.0 feature for token-based authentication.

Parameters

`[-enabled {true|false}]` - OAuth 2.0 Enabled

Use this parameter to enable or disable the OAuth 2.0 feature for the cluster.

Examples

The following example enables the OAuth 2.0 feature for the cluster:

```
cluster1::> security oauth2 modify -enabled true
```

security oauth2 show

Display global OAuth 2.0 configuration

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security oauth2 show` command displays the status of the OAuth 2.0 feature.

Examples

The following example displays the OAuth 2.0 feature status information -

```
cluster1::> security oauth2 show
                      Is OAuth 2.0 Enabled: true
```

security oauth2 client create

Configure OAuth 2.0 Provider

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security oauth2 client create` command creates OAuth 2.0 Provider configuration with the specified configuration name for token-based authentication. This command does not enable the OAuth 2.0 feature, it only configures it. Configuring and enabling the OAuth2.0 feature is a two-step process:

- Create an OAuth 2.0 Provider configuration using the `security oauth2 client create` command.

- Enable the OAuth 2.0 feature using the security oauth2 modify`-enabled`true` command. This step must be performed once regardless of the number of providers configured. If you have already enabled the OAuth 2.0 feature as part of configuring another provider, you do not have to perform this step again for this provider.

After an OAuth 2.0 configuration is created, it cannot be modified. It must be deleted and created again to change any settings.

> Enabling/Disabling OAuth 2.0 restarts the web server. Any HTTP/S connections that are active will be disrupted.

Parameters

`-config-name <text>` - Configuration Entry Name

This is the OAuth 2.0 configuration entry name.

`-application <OAuth 2.0 Applications>` - Application

This is the application for which OAuth 2.0 is configured. Currently only the *`http`* application is supported.

`-issuer {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - OAuth 2.0 Issuer

This is the OAuth 2.0 issuer to match with the "iss" field from the access token.

`[-audience <text>]` - OAuth 2.0 Audience

This is the OAuth 2.0 audience to match with the "aud" field from the access token. If this parameter is not set, then the "aud" field will not be matched and the REST API request will be forwarded to the provider with the matching "iss" field.

`[-client-id <text>]` - OAuth 2.0 Client ID

This is the Client identifier used in token introspection calls to the IdP server.

`[-introspection-endpoint {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}]` - OAuth 2.0 Token Introspection Endpoint Location

This is the URI of the desired IdP server used for token introspection.

`[-introspection-interval {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W | disabled}]` - OAuth 2.0 Token Introspection Refresh Interval in ISO-8601 format

This is the refresh interval in ISO-8601 format for caching the introspected tokens. When not set, the default value of *`0s`* is used, which caches introspected tokens for a period of "exp" value in the access token. This can be set to the value *`disabled`* to disable caching of tokens. Otherwise, it can be set to a value from *`1s`* to *`2147483647s`* .

`[-remote-user-claim <text>]` - OAuth 2.0 Remote User Claim

When the `-use-local-roles-if-present` parameter is set to true, and the token scope rules do not explicitly allow or deny the request, the value of the `-remote-user-claim` field will be used to find a match in the local user database for a user of the same name with the application of type *`http`* and authentication method *`password`* . When not set, the default value of *`sub`* is used.

`[-provider-jwks-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}]` - OAuth 2.0 Provider JSON Web Key Set Location

This is the URI where the JSON Web Key Set (JWKS) is hosted by the Identity Provider server.

[-jwks-refresh-interval {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W}] - OAuth 2.0 JSON Web Key Set Refresh Interval in ISO-8601 format

This is the refresh interval in ISO-8601 format for caching the JSON Web Key Set (JWKS) key set obtained from the provider-jwks-uri. When not set, the default value of *1h* is used. This can be set to a value from *3600s* to *2147483647s* .

[-outgoing-proxy <text>] - OAuth 2.0 Outgoing Proxy To Access External IdPs

This is the value for outgoing proxy to access external identity providers (IdPs). Use this parameter for local validation to download the JWKS key set in the JWKS URI, or for remote introspection for validating the access token in the Bearer field of the REST API request when the system is behind a proxy.

[-use-local-roles-if-present {true|false}] - Use Local Roles, If Present

When this parameter is set to *true* , and the scopes in the access token do not explicitly allow or deny the request, the local user role matching the user in the `-remote-user-claim` field (defaults to value of the "sub" field of the access token if not specified) will be checked for authorization of the request. The default value is *false* when not set, which means only the scopes in the access token is used to approve or deny the request.

[-skip-uri-validation {true|false}] - Skip URI Validations

When this parameter is set to *true* , validation of provider-jwks-uri is skipped. The default value of this parameter is *false* .

[-use-mutual-tls {none|request|required}] - Mutual TLS enforcement

This is the Mutual TLS setting for the OAuth 2.0 configuration. When set to *required* , OAuth 2.0 mutual TLS authentication is enforced for all access tokens and any token that does not have x5t#S256 property in the cnf section is rejected. The default value is *request* when not set, which means OAuth 2.0 mutual TLS authentication is enforced only if the x5t#S256 property is present in the cnf section of the access token. This can be disabled by setting to value *none* .

[-provider {basic|keycloak|auth0|adfs|entra}] - OAuth 2.0 Provider

This parameter optionally specifies the Identity Provider type. The default value of this parameter is *basic* .

Examples

The following example creates OAuth 2.0 Provider configuration for Local Validation:

```
cluster1::> security oauth2 client create -config-name auth1 -application
http -issuer https://issuer.example.com/ -provider-jwks-uri
https://issuer.example.com/.well-known/jwks.json -use-local-roles-if
-present true -remote-user-claim preferred_username -outgoing-proxy
https://outgoing_proxy
```

The following example creates OAuth 2.0 Provider configuration for Remote Introspection:

```
cluster1::> security oauth2 client create -config-name auth1 -application
http -issuer https://issuer.example.com/ -client-id client_id -client
-secret client_secret -use-local-roles-if-present true -remote-user-claim
preferred_username -outgoing-proxy https://outgoing_proxy -use-mutual-tls
required
```

Related Links

- security oauth2 modify

security oauth2 client delete

Delete OAuth 2.0 Provider

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security oauth2 client delete` command is used to remove an OAuth 2.0 Provider configuration.

Parameters

`-config-name <text>` - Configuration Entry Name
This is the OAuth 2.0 configuration entry name.

Examples

The following example removes the OAuth 2.0 Provider configuration named auth1:

```
cluster1::> security oauth2 client delete -config-name auth1
```

The following example removes all OAuth2.0 Provider configurations:

```
cluster1::> security oauth2 client delete -config-name *
```

security oauth2 client show

Display OAuth 2.0 Provider

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security oauth2 client show` command displays the configured OAuth 2.0 Provider configuration.

Parameters

{ [-fields <fieldname>,…]

This specifies the fields that need to be displayed.

| [-instance] }

If this parameter is specified, the command displays information about all OAuth 2.0 configuration entries.

[-config-name <text>] - Configuration Entry Name

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified config-name.

[-application <OAuth 2.0 Applications>] - Application

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified application. Currently only the *http* application is supported.

[-issuer {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}] - OAuth 2.0 Issuer

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified issuer.

[-audience <text>] - OAuth 2.0 Audience

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified audience.

[-client-id <text>] - OAuth 2.0 Client ID

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified client-id.

[-hashed-client-secret <Hex String>] - Hashed representation of client secret

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified hashed-client-secret.

[-introspection-endpoint {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}] - OAuth 2.0 Token Introspection Endpoint Location

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified introspection-endpoint.

[-introspection-interval {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W | disabled}] - OAuth 2.0 Token Introspection Refresh Interval in ISO-8601 format

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified introspection-interval.

[-remote-user-claim <text>] - OAuth 2.0 Remote User Claim

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified remote-user-claim.

[-provider-jwks-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}] - OAuth 2.0 Provider JSON Web Key Set Location

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified provider-jwks-uri.

[-jwks-refresh-interval {P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W}] - OAuth 2.0 JSON Web Key Set Refresh Interval in ISO-8601 format

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified jwks-refresh-interval.

[-outgoing-proxy <text>] - OAuth 2.0 Outgoing Proxy To Access External IdPs

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified outgoing-proxy.

[-use-local-roles-if-present {true|false}] - Use Local Roles, If Present

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified use-local-roles-if-present.

[-use-mutual-tls {none|request|required}] - Mutual TLS enforcement

This is the Mutual TLS setting for the OAuth 2.0 configuration. When set to *required* , OAuth 2.0 mutual TLS authentication is enforced for all access tokens and any token that does not have x5t#S256 property in the cnf section is rejected. The default value is *request* when not set, which means OAuth 2.0 mutual TLS authentication is enforced only if the x5t#S256 property is present in the cnf section of the access token. This can be disabled by setting to value *none* .

[-provider {basic|keycloak|auth0|adfs|entra}] - OAuth 2.0 Provider

If this parameter is specified, the command displays information only about the OAuth 2.0 configuration that match the specified Identity Provider.

Examples

The following example displays the OAuth 2.0 Provider configuration for Local Validation:

```
cluster1::> security oidc client show
                                Configuration Name: auth1
                                       Application: http
                           Issuer: https://issuer.example.com/
                                          Audience: -
                                         Client ID: -
                               Hashed Client Secret: -
                             Introspection Endpoint: -
                   Introspection Refresh Interval : -
                                    Use local roles: true
                 Provider JSON Web Key Set Location:
https://issuer.example.com/.well-known/jwks.json
                  JSON Web Key Set Refresh Interval: 1h
                                  Remote User Claim: preferred_username
                                     Outgoing Proxy: https://outgoing_proxy
                             Mutual TLS enforcement: request
                                           Provider: basic
```

The following example displays the OAuth 2.0 Provider configuration for Remote Introspection:

```
cluster1::> security oidc client show
                                 Configuration Name: auth1
                                        Application: http
                                             Issuer:
https://issuer.example.com/
                                           Audience: -
                                          Client ID: client_id
                                Hashed Client Secret:
e194e3472ee55c4202582cfbf59a03a37ef27085d2baf1b2fd7f7da3973c56fa
                              Introspection Endpoint: -
                     Introspection Refresh Interval : 0s
                                    Use local roles: true
                  Provider JSON Web Key Set Location: -
                   JSON Web Key Set Refresh Interval: -
                                  Remote User Claim: preferred_username
                                     Outgoing Proxy: https://outgoing_proxy
                              Mutual TLS enforcement: required
                                           Provider: basic
```

security oauth2 scope cli-to-scope generate

Generate OAuth 2.0 scope for the given CLI REST role creation command parameters

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security oauth2 scope cli-to-scope generate` command generates on ONTAP-specific OAuth 2.0 scope string based on local ONTAP custom roles created using security login rest-role create.

Parameters

`-role <text>` - Role name

The role name as in the security login rest-role create -role parameter. This parameter is required.

`-access <text>` - Access level

The access level as in the security login rest-role create -access parameter. Valid access levels are none, readonly, all, read_create, read_modify and read_create_modify. This parameter is required.

`[-api <text>]` - API path

The REST API URI as in the security login rest-role create -api parameter. Valid APIs start with /api/. This parameter is required.

`[-cluster-uuid <text>]` - Cluster UUID

The cluster UUID for which this scope applies. This parameter is optional. If not specified, the OAuth 2.0 scope is applicable to all clusters

Examples

To generate the OAuth 2.0 scope string applicable to all clusters for an ONTAP role named myrole for the REST API URI /api/cluster with admin (all) access:

```
cluster1::gt; security oauth2 scope cli-to-scope generate -role myrole
-api /api/cluster -access all -cluster-uuid *
ontap:*:myrole:all:*:/api/cluster
```

security oauth2 scope scope-to-cli generate

Generate CLI REST role command for the given OAuth 2.0 scope

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security oauth2 scope scope-to-cli generate` command generates on ONTAP CLI security login rest-role create command that is the equivalent of the specified OAuth 2.0 scope string.

Parameters

`-scopeString <text>` - OAuth 2.0 scope

The OAuth 2.0 scope string. This parameter is required.

Examples

To generate ONTAP CLI command given an OAuth 2.0 scope string applicable to all clusters for an ONTAP role named restclusterrole for the REST API URI /api/cluster with readonly access:

```
cluster1::gt; security oauth2 scope scope-to-cli generate -scopeString
ontap:*:restclusterrole:readonly:*:/api/cluster
Command for cluster <All>:
security login rest-role create -role restclusterrole -access readonly
-api /api/cluster
```

security protocol commands

security protocol modify

Modify application configuration options

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security protocol modify` command modifies the existing cluster-wide configuration of RSH and Telnet. Enable RSH and Telnet in the cluster by setting the enabled field as *`true`* .

Parameters

`-application <text>` - application (privilege: advanced)

Selects the application. Supported values are *`rsh`* and *`telnet`* .

`[-enabled {true|false}]` - enabled (privilege: advanced)

Enables or disables the corresponding application. The default value is *`false`* .

Examples

The following command enables RSH in the cluster. The default setting for RSH is *`false`* :

```
cluster1::> security protocol modify -application rsh -enabled true
```

The following command enables Telnet in the cluster. The default setting for Telnet is *`false`* :

```
cluster1::> security protocol modify -application telnet -enabled true
```

security protocol show

Show application configuration options

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security protocol show` command displays the cluster-wide configuration of RSH and Telnet in the cluster in advanced privilege mode. RSH and Telnet are disabled by default. Use the security protocol modify command to change the RSH and Telnet configuration that the cluster supports.

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.

`[-application <text>]` - application (privilege: advanced)

Displays the nonsecure applications in the cluster.

`[-enabled {true|false}]` - enabled (privilege: advanced)

Displays whether the application is enabled or disabled in the cluster.

Examples

The following example shows the default security protocol configurations for a cluster:

```
cluster1::> security protocol show

Application      Enabled
------------    -------------
rsh             false
telnet          false
```

The following example shows the security protocol configuration after RSH and Telnet have been enabled:

```
cluster1::> security protocol show
Application      Enabled
------------    -------------
rsh             true
telnet          true
```

Related Links

- security protocol modify

security protocol ssh modify

Modify the SSH configuration

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security protocol ssh modify` command modifies the existing cluster-wide configuration of SSH

Parameters

`[-per-source-limit <integer>]` - Per-Source Limit (privilege: advanced)

Modifies the maximum number of SSH instances per source IP address on a per-node basis.

`[-max-instances <integer>]` - Maximum Number of Instances (privilege: advanced)

Modifies the maximum number of SSH instances that can be handled on a per-node basis.

`[-connections-per-second <integer>]` - Connections Per Second (privilege: advanced)

Modifies the maximum number of SSH connections per second on a per-node basis.

Examples

The following example modifies cluster-wide SSH configuration:

```
cluster1::*> security protocol ssh modify -per-source-limit 30 -max
-instances 60 -connections-per-second 5
```

security protocol ssh show

Show the SSH configuration

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security protocol ssh show` command displays the cluster-wide SSH configuration in advanced privilege mode. Use the security protocol ssh modify command to change the SSH configuration that the cluster supports.

Examples

The following example displays cluster-wide SSH configuration:

```
cluster1::*> security protocol ssh show
Per-Source Limit: 32
Maximum Number of Instances: 64
     Connections Per Second: 10
```

Related Links

- security protocol ssh modify

security saml-sp commands

security saml-sp create

Configure SAML service provider for authentication

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security saml-sp create` command configures ONTAP with Security Assertion Markup Language (SAML) Service Provider (SP) for single sign-on authentication. This command does not enable SAML SP, it just configures it. Configuring and enabling SAML SP is a two-step process:

- Create a SAML SP configuration using `security saml-sp create` command.
- Enable SAML SP by using security saml-sp modify`-is-enabled`*true*

After the SAML SP configuration is created, it cannot be modified. It must be deleted and created again to change any settings.

This restarts the web server. Any HTTP/S connections that are active will be disrupted.

Parameters

`-idp-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')…}` - Identity Provider (IdP) Metadata Location

This is the URI of the desired identity provider's (IdP) metadata.

`[-sp-host <Remote InetAddress>]` - SAML Service Provider Host

This specifies the SAML service provider host IP address.

`{ -cert-ca <text>` - Server Certificate Issuing CA

This specifies the service provider's certificate issuing CA.

`-cert-serial <text>` - Server Certificate Serial Number

This specifies the service provider's certificate's serial number.

`| [-cert-common-name <FQDN or Custom Common Name>]` - Server Certificate Common Name `}`

This specifies the service provider certificate's common name.

`[-verify-metadata-server {true|false}]` - Verify IdP Metadata Server Identity

When the IdP metadata is downloaded, the identity of the server hosting the metadata is verified using Transport Layer Security (TLS), validating the server's X.509 certificate against the list of certificate authorities (CAs) in ONTAP, and verifying that the host in the server certificate matches the host in the URI (the `idp-uri` field). This verification can be bypassed by setting this field to *`false`* . Bypassing the server verification is not recommended as the server can not be trusted that way, but will be necessary to use non-TLS URIs, e.g. with the "http" scheme, or when the server certificates are self-signed. If the server's certificate was signed by a CA that is not installed in ONTAP, the security certificate install -type server-ca command can be used to install it.

`[-foreground {true|false}]` - Foreground Process

When this parameter is set to *`false`* the command runs in the background as a job. The default is *`true`* , which causes the command to return after the operation completes.

Examples

The following example configures ONTAP with SAML SP IdP information:

```
cluster1::> security saml-sp create -idp-uri http://public-idp-uri -sp
-host 1.1.1.1
    [Job 9] Job succeeded.
    cluster1::>
```

Related Links

- security saml-sp modify
- security certificate install

security saml-sp delete

Delete SAML service provider for authentication

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security saml-sp delete` command is used to remove the Security Access Markup Language (SAML) Service Provider (SP). Running this command frees resources used by the SP. SAML SP services will no longer be available after the SP is removed.

If the SAML SP is currently enabled, it is necessary to first use security saml-sp modify`-is-enabled`*false* `prior to `security saml-sp delete`. The security saml-sp modify`-is-enabled`*false* command must be issued by a password authenticated console application user or from a SAML authenticated command interface.

> This restarts the web server. Any HTTP/S connections that are active will be disrupted.

Examples

The following example unconfigures SAML SP:

```
cluster1::> security saml-sp delete
    cluster1::>
```

Related Links

- security saml-sp modify

security saml-sp modify

Modify SAML service provider authentication

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security saml-sp modify` command modifies the Security Assertion Markup Language (SAML) Service Provider (SP) configuration for single sign-on authentication. This command is used to enable or disable an existing SAML SP, `security saml-sp modify-is-enabled`*true* or *false*` respectively.

This command will check the validity of the current SAML SP configuration before enabling the SP. Also, it is necessary to use this command with the `-is-enabled`*false* `parameter prior to deleting an existing SAML SP configuration. SAML SP can only be disabled in this way by a password authenticated console application user or from a SAML authenticated command interface. The delete command must be used if the SAML configuration settings are to be changed, as only the `is-enabled` parameter can be modified.

> This may restart the web server. Any HTTP/S connections that are active may be disrupted.

Parameters

`[-is-enabled {true|false}]` - SAML Service Provider Enabled

Use this parameter to enable or disable the SAML SP.

Examples

The following example enables SAML SP:

```
cluster1::> security saml-sp modify -is-enabled true
    cluster1::>
```

security saml-sp repair

Repair a failed SAML SP configuation

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security saml-sp repair` command attempts to repair a failed SAML SP configuration on a given node. The status of the individual nodes can be viewed using the security saml-sp status show command.

This restarts the web server. Any active HTTP/S requests to the web server will be disrupted.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

This identifies a single node that matches the input. The repair job will run on this node.

`[-foreground {true|false}]` - Foreground Process (privilege: advanced)

When this parameter is set to *`false`* the command runs in the background as a job. The default is *`true`* , which causes the command to return after the operation completes.

Examples

The following example repairs a failed SAML SP configuration:

```
cluster1:> security saml-sp repair -node node-2
Warning: This restarts the web server. Any active HTTP/S requests to the
web
               server will be disrupted
Do you want to continue? {y|n}: y
      [Job 1321] Job succeeded.
cluster1:>
```

Related Links

- security saml-sp status show

security saml-sp show

Display SAML service provider for authentication

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security saml-sp show` command displays the Security Assertion Markup Language (SAML) Service Provider (SP) configuration.

The `Identity Provider (IdP) URI` indicates the URI of the desired IdP's metadata.

The `Service Provider (SP) host` indicates the IP address containing SAML SP metadata.

The `Certificate Common Name` indicates the SAML SP certificate's common name.

The `Certificate Serial` indicates the SAML SP certificate's serial number.

Examples

The following example displays the SAML SP configuration:

```
cluster1::> security saml-sp show
Identity Provider URI: https://www.my.idp.com
       Service Provider Host: 1.1.1.1
            Certificate Name: mycert
          Certificate Serial: 1234abcd
             Is SAML Enabled: false
```

security saml-sp status show

Display SAML service provider configuration status

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `security saml-sp status show` command displays the SAML Service Provider (SP) status for all nodes in the cluster.

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>|local}] - Node (privilege: advanced)

This identifies the node in the cluster.

[-status {not-configured|config-in-progress|config-failed|config-success}] - Update Status (privilege: advanced)

This identifies the SAML SP status on the specified node.

[-error-text <text>] - Error Text (privilege: advanced)

This identifies the error text associated with the latest saml SP update for this node.

[-is-enabled {true|false}] - SAML Service Provider Enabled (privilege: advanced)

When this parameter is set to *`true`* it indicates that the SAML SP is enabled on this node. Similarly, when this parameter is set to *`false`* , it indicates that the SAML SP is not enabled on this node.

Examples

The following example displays the SAML SP status information for all nodes in the cluster.

```
cluster::security saml-sp status> show
Node                             SAML SP Status         Enabled
-------------------------------  ---------------------  ----------
cluster-node1                    not-configured         false
cluster-node2                    not-configured         false
2 entries were displayed.

cluster::*>
```

security session commands

security session kill-cli

Kill a CLI session

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security session kill-cli` command is used to terminate CLI sessions. If the session being killed is actively processing a non-read command, the kill will wait until the command is complete before terminating the session. If the session being killed is actively processing a read (show) command, the kill will wait until the current row is returned before terminating the session.

Parameters

-node {<nodename>|local} - Node

Selects the sessions that match this parameter value. This identifies the node that is processing the session.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) that is processing the session.

[-start-time <MM/DD HH:MM:SS>] - Start Time

Selects the sessions that match this parameter value. This identifies the start time of the current active session.

-session-id <integer> - Session ID

Selects the sessions that match this parameter value. This number uniquely identifies a management session within a given node.

[-vserver <vserver>] - Vserver

Selects the sessions that match this parameter value. This identifies the Vserver associated with this management session.

[-username <text>] - Username

Selects the sessions that match this parameter value. This identifies the authenticated user associated with this management session.

[-application <text>] - Client Application

Selects the sessions that match this parameter value. This identifies the calling application by name.

[-location <text>] - Client Location

Selects the sessions that match this parameter value. This identifies the location of the calling client application. This is typically the IP address of the calling client, or "console" or "localhost" for console or localhost connections.

[-idle-seconds <integer>] - Idle Seconds

Selects the sessions that match this parameter value. When a session is not actively executing a command request (the session is idle), this indicates the time (in seconds) since the last request completed.

[-state {pending|active|idle}] - Session State

Selects the sessions that match this parameter value. This identifies the state (pending, active, or idle) of the session. The state is "pending" if it hit a session limit and the session is waiting for another session to end. The state is "idle" for CLI sessions that are waiting at the command prompt. The state is "active" if the session is actively working on a request.

[-request <text>] - Active Command

Selects the sessions that match this parameter value. This identifies the request (command) that is currently being handled by the session.

Examples

The following example illustrates killing a CLI session by specifying the node and the session id.

```
cluster1::> security session show -node node1

Node: node1               Interface: cli
Idle
Start Time     Sess ID Application Location           Vserver Username
Seconds
-------------- ------- ----------- ------------ ------------- --------
--------
03/27 16:58:13 1358    console     console           cluster1 admin
-
    Active Seconds: 0  Request: security session show
03/27 16:58:17 1359    ssh         10.98.16.164      cluster1 admin
650
2 entries were displayed.

cluster1::>

cluster1::> security session kill-cli -node node1 -session-id 1359
1 entry was acted on.

cluster1::> security session show -node node1

Node: node1               Interface: cli
Idle
Start Time     Sess ID Application Location           Vserver Username
Seconds
-------------- ------- ----------- ------------ ------------- --------
--------
03/27 16:58:13 1358    console     console           cluster1 admin
-
    Active Seconds: 0  Request: security session show

cluster1::>
```

The following example illustrates killing a CLI session by specifying the node and specifying a query on idle-seconds.

```
cluster1::> security session show -node node1

Node: node1               Interface: cli
Idle
Start Time     Sess ID Application Location           Vserver Username
Seconds
-------------- ------- ----------- ------------ ------------- --------
--------
03/27 16:58:13 1358    console     console           cluster1 admin
-
    Active Seconds: 0  Request: security session show
03/27 17:13:36 1479    ssh         10.98.16.164      cluster1 admin
83
2 entries were displayed.

cluster1::> security session kill-cli -node node1 -session-id * -idle
-seconds > 80
1 entry was acted on.

cluster1::> security session show

Node: node1               Interface: cli
Idle
Start Time     Sess ID Application Location           Vserver Username
Seconds
-------------- ------- ----------- ------------ ------------- --------
--------
03/27 16:58:13 1358    console     console           cluster1 admin
-
    Active Seconds: 0  Request: security session show

cluster1::>
```

security session show

Show current CLI, ONTAPI, and REST sessions

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security session show` command displays all active management sessions across the cluster.

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>|local}] - Node

Selects the sessions that match this parameter value. This identifies the node that is processing the session.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) that is processing the session.

[-start-time <MM/DD HH:MM:SS>] - Start Time

Selects the sessions that match this parameter value. This identifies the start time of the current active session.

[-session-id <integer>] - Session ID

Selects the sessions that match this parameter value. This number uniquely identifies a management session within a given node.

[-vserver <vserver>] - Vserver

Selects the sessions that match this parameter value. This identifies the Vserver associated with this management session.

[-username <text>] - Username

Selects the sessions that match this parameter value. This identifies the authenticated user associated with this management session.

[-application <text>] - Client Application

Selects the sessions that match this parameter value. This identifies the calling application by name.

[-location <text>] - Client Location

Selects the sessions that match this parameter value. This identifies the location of the calling client application. This is typically the IP address of the calling client, or "console" or "localhost" for console or localhost connections.

[-ipspace <IPspace>] - IPspace of Location

Selects the sessions that match this parameter value. This identifies the IPspace of the client location.

[-total <integer>] - Total Requests

Selects the sessions that match this parameter value. This identifies the total number of requests that have been made thus far in the active session. The following commands are not counted: top, up, cd, rows, history, exit.

[-failed <integer>] - Failed Requests

Selects the sessions that match this parameter value. This identifies the number of requests that have failed for any reason (including if they were blocked by configured limits).

[-max-time <integer>] - Maximum Time (ms)

Selects the sessions that match this parameter value. This identifies the maximum amount of time (in milliseconds) that any request took for this session.

[-last-time <integer>] - Last Time (ms)

Selects the sessions that match this parameter value. This identifies the amount of time (in milliseconds) that the last request took for this session.

[-total-seconds <integer>] - Total Seconds

Selects the sessions that match this parameter value. This identifies the total time (in seconds) that has been taken by all completed requests for the current session; it does not include session idle time.

[-state {pending|active|idle}] - Session State

Selects the sessions that match this parameter value. This identifies the state (pending, active, or idle) of the session. The state is "pending" if it hit a session limit and the session is waiting for another session to end. The state is "idle" for CLI sessions that are waiting at the command prompt. The state is "active" if the session is actively working on a request.

[-request <text>] - Request Input

Selects the sessions that match this parameter value. This identifies the request (command) that is currently being handled by the session.

[-idle-seconds <integer>] - Idle Seconds

Selects the sessions that match this parameter value. When a session is not actively executing a command request (the session is idle), this indicates the time (in seconds) since the last request completed.

[-active-seconds <integer>] - Active Seconds

Selects the sessions that match this parameter value. When a session is actively executing a command request, this indicates the time (in seconds) since the current request started.

Examples

The following example illustrates displaying all active sessions across the cluster. In this example, we see one active session on node *node2* from the *console* application. We also see three active sessions on node *node1* . One is from the *console* application and two are from the *ssh* application. Also one of the *ssh* sessions is from user *diag* and the other *ssh* session is from user *admin* .

```
cluster1::> security session show

Node: node1               Interface: cli
Idle
Start Time     Sess ID Application Location           Vserver Username
Seconds
-------------- ------- ----------- ------------ ------------- --------
--------
03/27 16:58:13 1358    console     console           cluster1 admin
-
    Active Seconds: 0  Request: security session show
03/27 17:17:04 1514    ssh         10.98.16.164      cluster1 admin
139
03/27 17:17:29 1515    ssh         10.98.16.164      cluster1 diag
115

Node: node2               Interface: cli
Idle
Start Time     Sess ID Application Location           Vserver Username
Seconds
-------------- ------- ----------- ------------ ------------- --------
--------
03/27 17:18:54 1509    console     console           cluster1 admin
23
4 entries were displayed.

cluster1::>
```

The following example illustrates displaying all active sessions that have been idle for longer than 500 seconds.

```
cluster1::> security session show -idle-seconds > 500

Node: node1               Interface: cli
Idle
Start Time     Sess ID Application Location           Vserver Username
Seconds
-------------- ------- ----------- ------------ ------------- --------
--------
03/27 17:17:04 1514    ssh         10.98.16.164      cluster1 admin
607
03/27 17:17:29 1515    ssh         10.98.16.164      cluster1 diag
583
2 entries were displayed.

cluster1::>
```

security session limit create

Create default session limit

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command allows creation of a default management session limit that does not yet exist. The default limits can be overridden for specific values within each category by using advanced privilege level commands.

Parameters

`-interface {cli|ontapi|rest}` - Interface

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-category {application|location|request|user|vserver}` - Category

The session type for this default limit. The following categories are supported: application, location, request, user, Vserver.

`-max-active-limit <integer>` - Max-Active Limit

The maximum number of concurrent sessions allowed for this interface and category.

Examples

The following example illustrates creating a default limit for management sessions using the same application.

```
cluster1::> security session limit create -interface ontapi -category
application -max-active-limit 8
```

security session limit delete

Delete default session limit

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command allows deletion of a default management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-category {application|location|request|user|vserver}` - Category

The session type for this default limit. The following categories are supported: application, location, request, user, Vserver.

Examples

The following example illustrates deleting all default limits for CLI management sessions.

```
cluster1::> security session limit delete -interface cli -category *
3 entries were deleted.
```

security session limit modify

Modify default session limit

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command allows modification of a default management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-category {application|location|request|user|vserver}` - Category

The session type for this default limit. The following categories are supported: application, location, request, user, Vserver.

`[-max-active-limit <integer>]` - Max-Active Limit

The maximum number of concurrent sessions allowed for this interface and category.

Examples

The following example illustrates modifying the default limit for CLI management sessions from the same location.

```
cluster1::> security session limit modify -interface cli -category
location -max-active-limit 4
```

security session limit show

Show default session limits

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command shows the default management session limits that have been configured for each interface and category.

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.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) to which the limit applies.

[-category {application|location|request|user|vserver}] - Category

Selects the sessions that match this parameter value. This identifies the category for the limit. The following categories are supported: application, location, request, user, and Vserver.

[-max-active-limit <integer>] - Max-Active Limit

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying the default limits for management sessions.

```
cluster1::> security session limit show
Interface Category    Max-Active
--------- ----------- ----------
cli       user                 2
cli       vserver              4
ontapi    vserver              2
3 entries were displayed.
```

security session limit application create

Create per-application session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows creation of a per-application management session limit that does not yet exist.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-application <text>` - Application (privilege: advanced)

The specified application to which this limit applies. The limit with the application name *`-default-`* is the limit used for any application without a specific configured limit.

`-max-active-limit <integer>` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and application.

Examples

The following example illustrates creating a limit for management sessions from a custom application.

```
cluster1::*> security session limit application create -interface ontapi
-application "custom_app" -max-active-limit 8
```

security session limit application delete

Delete per-application session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows deletion of a per-application management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-application <text>` - Application (privilege: advanced)

The specified application to which this limit applies. The limit with the application name *`-default-`* is the limit used for any application without a specific configured limit.

Examples

The following example illustrates deleting a limit for management sessions from a custom application.

```
cluster1::*> security session limit application delete -interface ontapi
-application "custom_app"
```

security session limit application modify

Modify per-application session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows modification of a per-application management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-application <text>` - Application (privilege: advanced)

The specified application to which this limit applies. The limit with the application name *`-default-`* is the limit used for any application without a specific configured limit.

`[-max-active-limit <integer>]` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and application.

Examples

The following example illustrates modifying management session limits for some custom applications.

```
cluster1::*> security session limit application modify -interface ontapi
-application custom* -max-active-limit 4
2 entries were modified.
```

security session limit application show

Show per-application session limits

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command shows the per-application management session limits that have been configured for each interface and application.

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.

[-interface {cli|ontapi|rest}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) to which the limit applies.

[-application <text>] - Application (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the application for the limit. The limit with the application name *`-default-`* is the limit used for any application without a specific configured limit.

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying the per-application limits for ONTAPI management sessions.

```
cluster1::*> security session limit application show -interface ontapi
Interface Application          Max-Active
--------- -------------------- ----------
ontapi    -default-                     5
ontapi    custom_app                   10
2 entries were displayed.
```

security session limit location create

Create per-location session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows creation of a per-location management session limit that does not yet exist.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-location <text>` - Location (privilege: advanced)

The specified location to which this limit applies. The limit with the location name *`-default-`* (in the *`Default`* IPspace) is the limit used for any location (in any IPspace) without a specific configured limit.

`[-ipspace <IPspace>]` - IPspace of Location (privilege: advanced)

This identifies the IPspace of the client location. If not specified, changes are made in the *`Default`* IPspace.

`-max-active-limit <integer>` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and location.

Examples

The following example illustrates creating a CLI limit for specific location.

```
cluster1::*> security session limit location create -interface cli
-location 10.98.16.164 -max-active-limit 1
```

security session limit location delete

Delete per-location session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows deletion of a per-location management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-location <text>` - Location (privilege: advanced)

The specified location to which this limit applies. The limit with the location name *`-default-`* (in the *`Default`* IPspace) is the limit used for any location (in any IPspace) without a specific configured limit.

`[-ipspace <IPspace>]` - IPspace of Location (privilege: advanced)

This identifies the IPspace of the client location. If not specified, changes are made in the *`Default`* IPspace.

Examples

The following example illustrates deleting limits for management sessions from a specific set of locations.

```
cluster1::*> security session limit location delete -interface * -location
10.98.*
3 entries were deleted.
```

security session limit location modify

Modify per-location session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows modification of a per-location management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-location <text>` - Location (privilege: advanced)

The specified location to which this limit applies. The limit with the location name *`-default-`* (in the *`Default`* IPspace) is the limit used for any location (in any IPspace) without a specific configured limit.

`[-ipspace <IPspace>]` - IPspace of Location (privilege: advanced)

This identifies the IPspace of the client location. If not specified, changes are made in the *`Default`* IPspace.

`[-max-active-limit <integer>]` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and location.

Examples

The following example illustrates modifying management sessions limits for specific locations.

```
cluster1::*> security session limit location modify -interface * -location
10.98.* -max-active-limit 2
3 entries were modified.
```

security session limit location show

Show per-location session limits

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command shows the per-location management session limits that have been configured for each interface and location.

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.

[-interface {cli|ontapi|rest}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) to which the limit applies.

[-location <text>] - Location (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the location for the limit. The limit with the location name *-default-* (only in the *Default* IPspace) is the limit used for any location (in any IPspace) without a specific configured limit.

[-ipspace <IPspace>] - IPspace of Location (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the IPspace of the client location. The default IPspace is *Default* .

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying the per-location limits for management sessions.

```
cluster1::*> security session limit location show
Interface Location             IPspace     Max-Active
--------- -------------------- ----------- ----------
cli       -default-            Default             16
cli       10.98.16.164         Default              0
ontapi    -default-            Default              6
ontapi    10.98.16.164         Default              0
4 entries were displayed.
```

security session limit request create

Create per-request session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows creation of a per-request management session limit that does not yet exist.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-request <text>` - Request Name (privilege: advanced)

The specified request to which this limit applies. The limit with the request name *`-default-`* is the limit used for any request without a specific configured limit.

`-max-active-limit <integer>` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and request.

Examples

The following example illustrates creating a limit for number of clients executing a specific API.

```
cluster1::*> security session limit request create -interface ontapi
-request storage-disk-get-iter -max-active-limit 2
```

security session limit request delete

Delete per-request session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows deletion of a per-request management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-request <text>` - Request Name (privilege: advanced)

The specified request to which this limit applies. The limit with the request name *`-default-`* is the limit used for any request without a specific configured limit.

Examples

The following example illustrates deleting custom limits for that were configured for the volume commands and APIs.

```
cluster1::*> security session limit request delete -interface * -request
volume*
4 entries were deleted.
```

security session limit request modify

Modify per-request session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows modification of a per-request management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-request <text>` - Request Name (privilege: advanced)

The specified request to which this limit applies. The limit with the request name *`-default-`* is the limit used for any request without a specific configured limit.

`[-max-active-limit <integer>]` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and request.

Examples

The following example illustrates modifying the limit of the number of clients simulatiously executing a specific API.

```
cluster1::*> security session limit request modify -interface ontapi
-request storage-disk-get-iter -max-active-limit 4
```

security session limit request show

Show per-request session limits

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command shows the per-request management session limits that have been configured for each interface and request.

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.

[-interface {cli|ontapi|rest}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) to which the limit applies.

[-request <text>] - Request Name (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the request (command or API) for the limit. The limit with the request name *-default-* is the limit used for any request without a specific configured limit.

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying the per-request limits for management sessions.

```
cluster1::*> security session limit request show
Interface Request                          Max-Active
--------- -------------------------------- ----------
cli       -default-                                10
ontapi    -default-                                 5
ontapi    storage-disk-get-iter                     2
3 entries were displayed.
```

security session limit user create

Create per-user session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows creation of a per-user management session limit that does not yet exist.

Parameters

-interface {cli|ontapi|rest} - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-vserver <vserver>` - Vserver (privilege: advanced)

The specified Vserver to which this limit applies. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

`-user <text>` - User (privilege: advanced)

The specified user to which this limit applies. The limit with the user name *`-default-`* is the limit used for any user without a specific configured limit.

`-max-active-limit <integer>` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface, Vserver, and user.

Examples

The following example illustrates creating a per-user limit override for ONTAPI requests for the *`admin`* user in the admin Vserver.

```
cluster1::*> security session limit user create -interface ontapi -vserver
cluster1 -username admin -max-active-limit 16
```

security session limit user delete

Delete per-user session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows deletion of a per-user management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-vserver <vserver>` - Vserver (privilege: advanced)

The specified Vserver to which this limit applies. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

`-user <text>` - User (privilege: advanced)

The specified user to which this limit applies. The limit with the user name *`-default-`* is the limit used for any user without a specific configured limit.

Examples

The following example illustrates deleting all user-specific limits for CLI management sessions.

```
cluster1::*> security session limit user delete -interface cli -user !"-
default-"
2 entries were deleted.
```

security session limit user modify

Modify per-user session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows modification of a per-user management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-vserver <vserver>` - Vserver (privilege: advanced)

The specified Vserver to which this limit applies. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

`-user <text>` - User (privilege: advanced)

The specified user to which this limit applies. The limit with the user name *`-default-`* is the limit used for any user without a specific configured limit.

`[-max-active-limit <integer>]` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface, Vserver, and user.

Examples

The following example illustrates modifying the admin user's limit for CLI management sessions.

```
cluster1::*> security session limit user modify -interface cli -vserver
cluster1 -username admin -max-active-limit 30
```

security session limit user show

Show per-user session limits

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command shows the per-user management session limits that have been configured for each interface, Vserver, and user.

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.

[-interface {cli|ontapi|rest}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) to which the limit applies.

[-vserver <vserver>] - Vserver (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the Vserver for the limit. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

[-user <text>] - User (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the user for the limit. The limit with the user name *`-default-`* is the limit used for any user without a specific configured limit.

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying the per-user limits for CLI management sessions. In this example, there is a default limit of 4 sessions for each user. That limit is expanded to 8 for the admin Vserver. That limit is further expanded to 20 for the *`admin`* user in the admin Vserver.

```
cluster1::*> security session limit user show -interface cli
Interface Vserver              User                Max-Active
--------- -------------------- ------------------- ----------
cli       Cluster              -default-                    4
cli       cluster1             -default-                    8
cli       cluster1             admin                       20
3 entries were displayed.
```

security session limit vserver create

Create per-vserver session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows creation of a per-Vserver management session limit that does not yet exist.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-vserver <vserver>` - Vserver (privilege: advanced)

The specified Vserver to which this limit applies. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

`-max-active-limit <integer>` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and Vserver.

Examples

The following example illustrates creating a per-Vserver limit override for ONTAPI requests on the admin Vserver.

```
cluster1::*> security session limit vserver create -interface ontapi
-vserver cluster1 -max-active-limit 4
```

security session limit vserver delete

Delete per-vserver session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows deletion of a per-Vserver management session limit. The "Cluster" vserver is used when the specific Vserver doesn’t have a configured limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-vserver <vserver>` - Vserver (privilege: advanced)

The specified Vserver to which this limit applies. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

Examples

The following example illustrates deleting all per-Vserver limits for management sessions except the default limit.

```
cluster1::*> security session limit vserver delete -interface * -vserver
!Cluster
1 entries was deleted.
```

security session limit vserver modify

Modify per-vserver session limit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command allows modification of a per-Vserver management session limit.

Parameters

`-interface {cli|ontapi|rest}` - Interface (privilege: advanced)

The interface (CLI, ONTAPI, or REST) to which the limit applies.

`-vserver <vserver>` - Vserver (privilege: advanced)

The specified Vserver to which this limit applies. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

`[-max-active-limit <integer>]` - Max-Active Limit (privilege: advanced)

The maximum number of concurrent sessions allowed for this interface and Vserver.

Examples

The following example illustrates modifying the admin Vserver's limit for CLI management sessions.

```
cluster1::*> security session limit vserver modify -interface cli -vserver
cluster1 -max-active-limit 40
```

security session limit vserver show

Show per-vserver session limits

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command shows the per-Vserver management session limits that have been configured for each interface and Vserver.

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.

[-interface {cli|ontapi|rest}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) to which the limit applies.

[-vserver <vserver>] - Vserver (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the Vserver for the limit. The "Cluster" Vserver is used to limit Vservers that do not have a configured limit.

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying the per-Vserver limits for management sessions.

```
cluster1::*> security session limit vserver show
Interface Vserver              Max-Active
--------- -------------------- ----------
cli       Cluster                       4
ontapi    Cluster                       2
ontapi    cluster1                     16
3 entries were displayed.
```

security session request-statistics show-by-application

Show session request statistics by application

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security session request-statistics show-by-application` command shows historical statistics for management session activity, categorized by application name. CLI sessions connections will have an application name based on the connection method, i.e.: *ssh* , *telnet* , *rsh* , *console* , or *ngsh* . ONTAPI sessions will extract the application name from the ZAPI request. ONTAP looks for the application name in the following three locations, in the following order of precedence:

```
     1. The "X-Dot-Client-App" HTTP header;
     2. The "app-name" attribute of the "netapp" element, within the ZAPI
XML request;
     3. The "User-Agent" HTTP header.
```

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>|local}] - Node

Selects the sessions that match this parameter value. This identifies the node that processed the session.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) that processed the session.

[-application <text>] - Application

Selects the sessions that match this parameter value. This identifies the calling application by name.

[-total <integer>] - Total Requests

Selects the sessions that match this parameter value. This identifies the total number of requests that have been made on a session. The following commands are not counted: top, up, cd, rows, history, exit.

[-blocked <integer>] - Blocked Requests

Selects the sessions that match this parameter value. This identifies the number of requests that were blocked due to configured limits.

[-failed <integer>] - Failed Requests

Selects the sessions that match this parameter value. This identifies the number of requests that failed for any reason (including if they were blocked by configured limits).

[-max-time <integer>] - Maximum Time (ms)

Selects the sessions that match this parameter value. This identifies the maximum amount of time (in milliseconds) that any request took.

[-last-time <integer>] - Last Time (ms)

Selects the sessions that match this parameter value. This identifies the amount of time (in milliseconds) that the last request took.

[-active <integer>] - Number Active Now

Selects the sessions that match this parameter value. This identifies the number of currently active sessions.

[-max-active <integer>] - Max Number Active

Selects the sessions that match this parameter value. This identifies the maximum number of concurrently active sessions.

[-last-active-seconds <integer>] - Seconds Since Last Session Start

Selects the sessions that match this parameter value. When a session is active, this indicates the time (in seconds) since the last session started.

[-idle-seconds <integer>] - Idle Seconds

Selects the sessions that match this parameter value. When no sessions are active, this indicates the time (in seconds) since the last session ended.

[-total-seconds <integer>] - Total Seconds

Selects the sessions that match this parameter value. This identifies the total time (in seconds) that have been taken by all completed requests; it does not include session idle time.

[-average-time <integer>] - Average Time (ms)

Selects the sessions that match this parameter value. This identifies the mean time spent processing requests.

[-success-percent <percent>] - Success Percent

Selects the sessions that match this parameter value. This identifies the percentage of successful requests.

[-blocked-percent <percent>] - Blocked Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that were blocked due to configured limits.

[-failed-percent <percent>] - Failed Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that failed for any reason (including if they were blocked by configured limits).

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying historical statistics for all management session activity across the cluster, categorized by application name.

```
cluster1::> security session request-statistics show-by-application

Node: node1               Interface: cli                 Idle    Total
Application                  Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
console                       2126   0   6  95%   96       68      361
170
ssh                              6   2   3 100%    0        -      794
132444

Node: node1               Interface: ontapi              Idle    Total
Application                  Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
api_test                         2   0   1 100%    0       13        0
18

Node: node2               Interface: cli                 Idle    Total
Application                  Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
console                       2090   0   6  95%   96       90      655
313
4 entries were displayed.

cluster1::>
```

The following example illustrates displaying historical statistics for management session activity on a specific node and for a specific application.

```
cluster1::> security session request-statistics show-by-application -node
node1 -application api_test

Node: node1                Interface: ontapi              Idle    Total
Application                   Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
api_test                          2   0   1 100%    0      102        0
18

cluster1::>
```

security session request-statistics show-by-location

Show session request statistics by location

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security session request-statistics show-by-location` command shows historical statistics for management session activity, categorized by client location.

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>|local}] - Node

Selects the sessions that match this parameter value. This identifies the node that processed the session.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) that processed the session.

[-location <text>] - Client Location

Selects the sessions that match this parameter value. This identifies the location of the calling client application. This is typically the IP address of the calling client, or "console" or "localhost" for console or localhost connections.

[-ipspace <IPspace>] - IPspace of Location

Selects the sessions that match this parameter value. This identifies the IPspace of the client location.

[-total <integer>] - Total Requests

Selects the sessions that match this parameter value. This identifies the total number of requests that have been made on a session. The following commands are not counted: top, up, cd, rows, history, exit.

[-blocked <integer>] - Blocked Requests

Selects the sessions that match this parameter value. This identifies the number of requests that were blocked due to configured limits.

[-failed <integer>] - Failed Requests

Selects the sessions that match this parameter value. This identifies the number of requests that failed for any reason (including if they were blocked by configured limits).

[-max-time <integer>] - Maximum Time (ms)

Selects the sessions that match this parameter value. This identifies the maximum amount of time (in milliseconds) that any request took.

[-last-time <integer>] - Last Time (ms)

Selects the sessions that match this parameter value. This identifies the amount of time (in milliseconds) that the last request took.

[-active <integer>] - Number Active Now

Selects the sessions that match this parameter value. This identifies the number of currently active sessions.

[-max-active <integer>] - Max Number Active

Selects the sessions that match this parameter value. This identifies the maximum number of concurrently active sessions.

[-last-active-seconds <integer>] - Seconds Since Last Session Start

Selects the sessions that match this parameter value. When a session is active, this indicates the time (in seconds) since the last session started.

[-idle-seconds <integer>] - Idle Seconds

Selects the sessions that match this parameter value. When no sessions are active, this indicates the time (in seconds) since the last session ended.

[-total-seconds <integer>] - Total Seconds

Selects the sessions that match this parameter value. This identifies the total time (in seconds) that have been taken by all completed requests; it does not include session idle time.

[-average-time <integer>] - Average Time (ms)

Selects the sessions that match this parameter value. This identifies the mean time spent processing requests.

[-success-percent <percent>] - Success Percent

Selects the sessions that match this parameter value. This identifies the percentage of successful requests.

[-blocked-percent <percent>] - Blocked Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that were blocked due to configured limits.

[-failed-percent <percent>] - Failed Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that failed for any reason (including if they were blocked by configured limits).

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying historical statistics for all management session activity across the cluster, categorized by location.

```
cluster1::> security session request-statistics show-by-location

Node: node1                Interface: cli                  Idle    Total
Location          IPspace    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
----------------- ------- -------- --- --- ---- ---- -------- --------
--------
console           Default       21   1   1 100%    0        -      127
6063
localhost         Default     2523   0   5  95%  115       20      280
111

Node: node1                Interface: ontapi               Idle    Total
Location          IPspace    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
----------------- ------- -------- --- --- ---- ---- -------- --------
--------
10.98.17.254      Default        2   0   1 100%    0     2419        0
18

Node: node2                Interface: cli                  Idle    Total
Location          IPspace    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
----------------- ------- -------- --- --- ---- ---- -------- --------
--------
console           Default        6   0   1  83%    1     2941      423
70557
localhost         Default     2502   0   5  95%  114       41      277
110
7 entries were displayed.

cluster1::>
```

The following example illustrates displaying historical statistics for management session activity on a specific

node and for a specific location.

```
cluster1::> security session request-statistics show-by-location -node
node2 -location localhost

Node: node2               Interface: cli                 Idle    Total
Location          IPspace    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
----------------- ------- -------- --- --- ---- ---- -------- --------
--------
localhost         Default     2524   0   5  95%  115       30      279
110

cluster1::>
```

security session request-statistics show-by-request

Show session request statistics by request name

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security session request-statistics show-by-request` command shows historical statistics for management session activity, categorized by request (command or API name).

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>|local}] - Node

Selects the sessions that match this parameter value. This identifies the node that processed the session.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) that processed the session.

[-request <text>] - Request Name

Selects the sessions that match this parameter value. This identifies the command associated with these requests.

[-total <integer>] - Total Requests

Selects the sessions that match this parameter value. This identifies the total number of requests that have

been made on a session. The following commands are not counted: top, up, cd, rows, history, exit.

[-blocked <integer>] - Blocked Requests

Selects the sessions that match this parameter value. This identifies the number of requests that were blocked due to configured limits.

[-failed <integer>] - Failed Requests

Selects the sessions that match this parameter value. This identifies the number of requests that failed for any reason (including if they were blocked by configured limits).

[-max-time <integer>] - Maximum Time (ms)

Selects the sessions that match this parameter value. This identifies the maximum amount of time (in milliseconds) that any request took.

[-last-time <integer>] - Last Time (ms)

Selects the sessions that match this parameter value. This identifies the amount of time (in milliseconds) that the last request took.

[-active <integer>] - Number Active Now

Selects the sessions that match this parameter value. This identifies the number of currently active requests.

[-max-active <integer>] - Max Number Active

Selects the sessions that match this parameter value. This identifies the maximum number of concurrently active requests.

[-last-active-seconds <integer>] - Seconds Since Last Request Start

Selects the sessions that match this parameter value. When requests are active, this indicates the time (in seconds) since the last request started.

[-idle-seconds <integer>] - Idle Seconds

Selects the sessions that match this parameter value. When no requests are active, this indicates the time (in seconds) since the last request ended.

[-total-seconds <integer>] - Total Seconds

Selects the sessions that match this parameter value. This identifies the total time (in seconds) that have been taken by all completed requests; it does not include session idle time.

[-average-time <integer>] - Average Time (ms)

Selects the sessions that match this parameter value. This identifies the mean time spent processing requests.

[-success-percent <percent>] - Success Percent

Selects the sessions that match this parameter value. This identifies the percentage of successful requests.

[-blocked-percent <percent>] - Blocked Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that were blocked due to configured limits.

[-failed-percent <percent>] - Failed Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that failed for any reason (including if they were blocked by configured limits).

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying historical statistics for all management session activity on a specific node, with a specific request query.

```
cluster1::> security session request-statistics show-by-request -node
node1 -request network*

Node: node1                Interface: cli                  Idle    Total
Request Name                  Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
network interface create         2   0   1 100%    0     2556        0
485
network interface modify         1   0   1 100%    0     2518        0
34
network interface show           8   0   1 100%    0     2152       12
1614
network route create             1   0   1 100%    0     2135        0
45
network route show               2   0   1 100%    0     2145        0
17
5 entries were displayed.

cluster1::>
```

security session request-statistics show-by-user

Show session request statistics by username

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security session request-statistics show-by-user` command shows historical statistics for management session activity, categorized by username. Entries for username 'autosupport' reflect commands that are executed by the AutoSupport OnDemand feature.

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>|local}] - Node

Selects the sessions that match this parameter value. This identifies the node that processed the session.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) that processed the session.

[-vserver <vserver>] - Vserver

Selects the sessions that match this parameter value. This identifies the Vserver associated with this management session.

[-username <text>] - Username

Selects the sessions that match this parameter value. This identifies the authenticated user associated with this management session.

[-total <integer>] - Total Requests

Selects the sessions that match this parameter value. This identifies the total number of requests that have been made on a session. The following commands are not counted: top, up, cd, rows, history, exit.

[-blocked <integer>] - Blocked Requests

Selects the sessions that match this parameter value. This identifies the number of requests that were blocked due to configured limits.

[-failed <integer>] - Failed Requests

Selects the sessions that match this parameter value. This identifies the number of requests that failed for any reason (including if they were blocked by configured limits).

[-max-time <integer>] - Maximum Time (ms)

Selects the sessions that match this parameter value. This identifies the maximum amount of time (in milliseconds) that any request took.

[-last-time <integer>] - Last Time (ms)

Selects the sessions that match this parameter value. This identifies the amount of time (in milliseconds) that the last request took.

[-active <integer>] - Number Active Now

Selects the sessions that match this parameter value. This identifies the number of currently active sessions.

[-max-active <integer>] - Max Number Active

Selects the sessions that match this parameter value. This identifies the maximum number of concurrently

active sessions.

[-last-active-seconds <integer>] - Seconds Since Last Session Start

Selects the sessions that match this parameter value. When a session is active, this indicates the time (in seconds) since the last session started.

[-idle-seconds <integer>] - Idle Seconds

Selects the sessions that match this parameter value. When no sessions are active, this indicates the time (in seconds) since the last session ended.

[-total-seconds <integer>] - Total Seconds

Selects the sessions that match this parameter value. This identifies the total time (in seconds) that have been taken by all completed requests; it does not include session idle time.

[-average-time <integer>] - Average Time (ms)

Selects the sessions that match this parameter value. This identifies the mean time spent processing requests.

[-success-percent <percent>] - Success Percent

Selects the sessions that match this parameter value. This identifies the percentage of successful requests.

[-blocked-percent <percent>] - Blocked Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that were blocked due to configured limits.

[-failed-percent <percent>] - Failed Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that failed for any reason (including if they were blocked by configured limits).

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying historical statistics for all management session activity across the cluster, categorized by username.

```
cluster1::> security session request-statistics show-by-user

Node: node1                Interface: cli                  Idle    Total
Vserver          Username    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
-------------- ---------- -------- --- --- ---- ---- -------- --------
--------
cluster1            admin       81   1   3  80%   16        -     1228
15171
                     diag        1   0   1 100%    0     1982     1511
1511958
              autosupport        4   0   1 100%    0        -        0
17

Node: node1                Interface: ontapi               Idle    Total
Vserver          Username    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
-------------- ---------- -------- --- --- ---- ---- -------- --------
--------
cluster1            admin        2   0   1 100%    0     2585        0
18

Node: node2                Interface: cli                  Idle    Total
Vserver          Username    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
-------------- ---------- -------- --- --- ---- ---- -------- --------
--------
cluster1            admin        6   1   1  83%    1     3106      423
70557
4 entries were displayed.

cluster1::>
```

The following example illustrates displaying historical statistics for management session activity on a specific node and for a specific username.

```
cluster1::> security session request-statistics show-by-user -node node1
-username diag

Node: node1                Interface: cli                  Idle    Total
Vserver          Username    Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
-------------- ---------- -------- --- --- ---- ---- -------- --------
--------
cluster1             diag        1   0   1 100%    0        -     1511
1511958

cluster1::>
```

security session request-statistics show-by-vserver

Show session request statistics by Vserver

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `security session request-statistics show-by-vserver` command shows historical statistics for management session activity, categorized by vserver.

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>|local}] - Node

Selects the sessions that match this parameter value. This identifies the node that processed the session.

[-interface {cli|ontapi|rest}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI, ONTAPI, or REST) that processed the session.

[-vserver <vserver>] - Vserver

Selects the sessions that match this parameter value. This identifies the Vserver associated with this management session.

[-total <integer>] - Total Requests

Selects the sessions that match this parameter value. This identifies the total number of requests that have been made on a session. The following commands are not counted: top, up, cd, rows, history, exit.

[-blocked <integer>] - Blocked Requests

Selects the sessions that match this parameter value. This identifies the number of requests that were blocked due to configured limits.

[-failed <integer>] - Failed Requests

Selects the sessions that match this parameter value. This identifies the number of requests that failed for any reason (including if they were blocked by configured limits).

[-max-time <integer>] - Maximum Time (ms)

Selects the sessions that match this parameter value. This identifies the maximum amount of time (in milliseconds) that any request took.

[-last-time <integer>] - Last Time (ms)

Selects the sessions that match this parameter value. This identifies the amount of time (in milliseconds) that the last request took.

[-active <integer>] - Number Active Now

Selects the sessions that match this parameter value. This identifies the number of currently active sessions.

[-max-active <integer>] - Max Number Active

Selects the sessions that match this parameter value. This identifies the maximum number of concurrently active sessions.

[-last-active-seconds <integer>] - Seconds Since Last Session Start

Selects the sessions that match this parameter value. When a session is active, this indicates the time (in seconds) since the last session started.

[-idle-seconds <integer>] - Idle Seconds

Selects the sessions that match this parameter value. When no sessions are active, this indicates the time (in seconds) since the last session ended.

[-total-seconds <integer>] - Total Seconds

Selects the sessions that match this parameter value. This identifies the total time (in seconds) that have been taken by all completed requests; it does not include session idle time.

[-average-time <integer>] - Average Time (ms)

Selects the sessions that match this parameter value. This identifies the mean time spent processing requests.

[-success-percent <percent>] - Success Percent

Selects the sessions that match this parameter value. This identifies the percentage of successful requests.

[-blocked-percent <percent>] - Blocked Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that were blocked due to configured limits.

[-failed-percent <percent>] - Failed Percent

Selects the sessions that match this parameter value. This identifies the percentage of requests that failed for any reason (including if they were blocked by configured limits).

[-max-active-limit <integer>] - Max-Active Limit (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the configured limit that is used to throttle or reject requests.

Examples

The following example illustrates displaying historical statistics for all management session activity across the cluster, categorized by Vserver.

```
cluster1::> security session request-statistics show-by-vserver

Node: node1                Interface: cli                 Idle    Total
Vserver                       Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
cluster1                       2725   1   8  94%  146        -     3052
1120

Node: node1                Interface: ontapi              Idle    Total
Vserver                       Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
cluster1                          2   0   1 100%    0     2742        0
18

Node: node2                Interface: cli                 Idle    Total
Vserver                       Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
cluster1                       2552   1   6  95%  117        -      705
276
3 entries were displayed.

cluster1::>
```

The following example illustrates displaying historical statistics for management session activity on a specific node, for a specific Vserver.

```
cluster1::> security session request-statistics show-by-vserver -node
node1 -vserver cluster1

Node: node1                Interface: cli                  Idle    Total
Vserver                       Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
cluster1                       2747   1   8  94%  147        -     3055
1112

Node: node1                Interface: ontapi               Idle    Total
Vserver                       Total Now Max Pass Fail  Seconds  Seconds Avg
(ms)
------------------------- -------- --- --- ---- ---- -------- --------
--------
cluster1                           2   0   1 100%    0     2902        0
18
2 entries were displayed.

cluster1::>
```

security ssh commands

security ssh add

Add SSH configuration options

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The ` security ssh add ` command adds additional SSH key exchange algorithms or ciphers or MAC algorithms to the existing configurations of the cluster or a Vserver. The added algorithms or ciphers or MAC algorithms are enabled on the cluster or Vserver. If you change the cluster configuration settings, it is used as the default for all newly created Vservers. The existing SSH key exchange algorithms, ciphers, and MAC algorithms remain unchanged in the configuration. If the SSH key exchange algorithms or ciphers or MAC algorithms are already enabled in the current configuration, the command will does not not fail. ONTAP supports the *diffie-hellman-group-exchange-sha256* , *diffie-hellman-group16-sha512* and *diffie-hellman-group18-sha512* key exchange algorithms for SHA-2. ONTAP also supports the *diffie-hellman-group-exchange-sha1* , *diffie-hellman-group14-sha1* , and *diffie-hellman-group1-sha1* SSH key exchange algorithms for SHA-1. The SHA-2 key exchange algorithm is more secure than the SHA-1 key exchange algorithms. ONTAP also supports *ecdh-sha2-nistp256* , *ecdh-sha2-nistp384* , *ecdh-sha2-nistp521* , and *curve25519-sha256* . ONTAP also supports the AES and 3DES symmetric encryptions (also known as ciphers) of the following types: *aes256-ctr* , *aes192-ctr* , *aes128-ctr* , *aes256-cbc* , *aes192-cbc* , *aes128-cbc* , *aes128-gcm* , *aes256-gcm* , and *3des-cbc* . ONTAP supports MAC algorithms of the following types: *hmac-sha1* , *hmac-sha1-96* , *hmac-md5* , *hmac-md5-96* , *umac-64* ,

umac-64 , *umac-128* , *hmac-sha2-256* , *hmac-sha2-512* , *hmac-sha1-etm* , *hmac-sha1-96-etm* , *hmac-sha2-256-etm* , *hmac-sha2-512-etm* , *hmac-md5-etm* , *hmac-md5-96-etm* , *umac-64-etm* , and *umac-128-etm* .

Parameters

`-vserver <Vserver Name>` - Vserver

Identifies the Vserver to which you want to add additional SSH key exchange algorithms or ciphers.

`[-key-exchange-algorithms <algorithm name>,…]` - List of SSH Key Exchange Algorithms to Add

Adds the specified SSH key exchange algorithm or algorithms to the Vserver.

`[-ciphers <cipher name>,…]` - List of SSH Ciphers to Add

Adds the specified cipher or ciphers to the Vserver.

`[-mac-algorithms <MAC name>,…]` - List of SSH MAC Algorithms to Add

Adds the specified MAC algorithm or algorithms to the Vserver.

`[-host-key-algorithms <HostKey Algorithms>,…]` - List of SSH Host Key Algorithms to Add

Adds the specified host key algorithms to the Vserver.

Examples

The following command adds the *diffie-hellman-group-exchange-sha256* and *diffie-hellman-group-exchange-sha1* key exchange algorithms for the cluster1 Vserver. It also adds the *aes256-cbc* and *aes192-cbc* ciphers and the *hmac-sha1* and *hmac-sha2-256* MAC algorithms to the cluster1 Vserver.

```
cluster1::> security ssh add -vserver cluster1 -key-exchange-algorithms
diffie-hellman-group-exchange-sha256,diffie-hellman-group-exchange-sha1
-ciphers aes256-cbc,aes192-cbc -mac-algorithms hmac-sha1,hmac-sha2-256
-host-key-algorithms ecdsa-sha2-nistp256,ssh-rsa
```

security ssh modify

Modify SSH configuration options

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security ssh modify` command replaces the existing configurations of the SSH key exchange algorithms, ciphers, MAC algorithms, maximum authentication retry count, host key algorithms and whether *ssh-rsa* signature scheme is enabled for RSA keys in publickey algorithms, for the cluster or a Vserver, with the configuration settings you specify. If you modify the cluster configuration settings, it will be used as the default for all newly created Vservers. ONTAP supports the *diffie-hellman-group-exchange-sha256* , *diffie-hellman-group16-sha512* and *diffie-hellman-group18-sha512* key exchange algorithms for SHA-2. ONTAP also supports the *diffie-hellman-group-exchange-sha1* , *diffie-hellman-group14-sha1* and *diffie-hellman-group1-sha1* SSH key exchange algorithms for SHA-1. The SHA-2

key exchange algorithm is more secure than the SHA-1 key exchange algorithms. ONTAP also supports the AES and 3DES symmetric encryptions (also known as ciphers) of the following types: *aes256-ctr* , *aes192-ctr* , *aes128-ctr* , *aes256-cbc* , *aes192-cbc* , *aes128-cbc* , *aes128-gcm* , *aes256-gcm* and *3des-cbc* . ONTAP supports MAC algorithms of the following types: *hmac-sha1* , *hmac-sha1-96* , *hmac-md5* , *hmac-md5-96* , *umac-64* , *umac-64* , *umac-128* , *hmac-sha2-256* , *hmac-sha2-512* , *hmac-sha1-etm* , *hmac-sha1-96-etm* , *hmac-sha2-256-etm* , *hmac-sha2-512-etm* , *hmac-md5-etm* , *hmac-md5-96-etm* , *umac-64-etm* and *umac-128-etm* . ONTAP supports host key algorithms of the following types: *ecdsa-sha2-nistp256* , *ssh-rsa* and *ssh-ed25519* .

Parameters

-vserver <Vserver Name> - Vserver

Identifies the Vserver for which you want to replace the existing SSH configurations.

[-key-exchange-algorithms <algorithm name>,…] - Key Exchange Algorithms

Enables the specified SSH key exchange algorithm or algorithms for the Vserver. This parameter also replaces all existing SSH key exchange algorithms with the specified settings.

[-ciphers <cipher name>,…] - Ciphers

Enables the specified cipher or ciphers for the Vserver. This parameter also replaces all existing ciphers with the specified settings.

[-mac-algorithms <MAC name>,…] - MAC Algorithms

Enables the specified MAC algorithm or algorithms for the Vserver. This parameter also replaces all existing MAC algorithms with the specified settings.

[-max-authentication-retry-count <integer>] - Max Authentication Retry Count

Modifies the maximum number of authentication retry count for the Vserver.

[-host-key-algorithms <HostKey Algorithms>,…] - Host Key Algorithms

Enables the specified host key algorithm or algorithms for the Vserver. This parameter also replaces all existing host key algorithms with the specified settings.

[-is-rsa-in-publickey-algorithms-enabled {true|false}] - Is ssh-rsa in Publickey Algorithms Enabled

Modifies the flag to enable or disable *ssh-rsa* signature scheme, which is based on SHA-1 hash algorithm, for RSA keys in publickey algorithms.

Examples

The following command enables the *diffie-hellman-group-exchange-sha256* and *diffie-hellman-group14-sha1* key exchange algorithms for the "cluster1" Vserver. It also enables the *aes256-ctr* , *aes192-ctr* and *aes128-ctr* ciphers, *hmac-sha1* and *hmac-sha2-256* MAC algorithms for the "cluster1" Vserver. It also disables *ssh-rsa* signature scheme for RSA keys in publickey algorithms and modifies the maximum authentication retry count to 3 for the "cluster1" Vserver:

```
cluster1::> security ssh modify -vserver cluster1 -key-exchange-algorithms
diffie-hellman-group-exchange-sha256,diffie-hellman-group14-sha1 -ciphers
aes256-ctr,aes192-ctr,aes128-ctr -mac-algorithms hmac-sha1,hmac-sha2-256
-is-rsa-in-publickey-algorithms-enabled false -max-authentication-retry
-count 3
```

security ssh prepare-to-downgrade

Downgrade the SSH configuration to be compatible with releases earlier than Data ONTAP 9.2.0.

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command downgrades the SSH configurations of all Vservers and the cluster to settings compatible with releases earlier than ONTAP 9.2.0. This command also disables the max-authentication-retry feature. You must run this command in advanced privilege mode when prompted to do so during the release downgrade. Otherwise, the release downgrade process will fail.

Examples

The following command downgrades the SSH security configurations of all Vservers and the cluster to settings compatible with releases earlier than ONTAP 9.2.0.

```
cluster1::*> security ssh prepare-to-downgrade
```

security ssh remove

Remove SSH configuration options

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The ` security ssh remove` command removes the specified SSH key exchange algorithms or ciphers from the existing configurations of the cluster or a Vserver. The removed algorithms or ciphers are disabled on the cluster or Vserver. If you changed the cluster configuration settings, it will be used as the default for all newly created Vservers. If the SSH key exchange algorithms or ciphers that you specify with this command are not currently enabled, the command does not fail. ONTAP supports the *diffie-hellman-group-exchange-sha256* , *diffie-hellman-group16-sha512* and *diffie-hellman-group18-sha512* key exchange algorithms for SHA-2. ONTAP also supports the *diffie-hellman-group-exchange-sha1* , *diffie-hellman-group14-sha1* , and *diffie-hellman-group1-sha1* SSH key exchange algorithms for SHA-1. The SHA-2 key exchange algorithm is more secure than the SHA-1 key exchange algorithms. ONTAP also supports *ecdh-sha2-nistp256* , *ecdh-sha2-nistp384* , *ecdh-sha2-nistp521* , and *curve25519-sha256* . ONTAP also supports the AES and 3DES symmetric encryption (also known as ciphers) of the following types: *aes256-ctr* , *aes192-ctr* , *aes128-ctr* , *aes256-cbc* , *aes192-cbc* , *aes128-cbc* , *aes128-gcm* , *aes256-gcm* and *3des-cbc* . ONTAP supports MAC algorithms of the following types: *hmac-*

sha1 , *hmac-sha1-96* , *hmac-md5* , *hmac-md5-96* , *umac-64* , *umac-64* , *umac-128* , *hmac-sha2-256* , *hmac-sha2-512* , *hmac-sha1-etm* , *hmac-sha1-96-etm* , *hmac-sha2-256-etm* , *hmac-sha2-512-etm* , *hmac-md5-etm* , *hmac-md5-96-etm* , *umac-64-etm* , and *umac-128-etm* .

Parameters

`-vserver <Vserver Name>` - Vserver

Identifies the Vserver from which you want to remove the SSH key exchange algorithms or ciphers.

`[-key-exchange-algorithms <algorithm name>,…]` - List of SSH Key Exchange Algorithms to Remove

Removes the specified key exchange algorithm or algorithms from the Vserver.

`[-ciphers <cipher name>,…]` - List of SSH Ciphers to Remove

Removes the specified cipher or ciphers from the Vserver.

`[-mac-algorithms <MAC name>,…]` - List of SSH MAC Algorithms to Remove

Removes the specified MAC algorithm or algorithms from the Vserver.

`[-host-key-algorithms <HostKey Algorithms>,…]` - List of SSH Host Key Algorithms to Remove

Removes the specified host key algorithms from the Vserver.

Examples

The following command removes the *diffie-hellman-group1-sha1* and *diffie-hellman-group-exchange-sha1* key exchange algorithms from the cluster1 Vserver. It also removes the *aes128-cbc* and *3des-cbc* ciphers and the *hmac-sha1-96* and *hmac-sha2-256* MAC algorithms from the cluster1 Vserver.

```
cluster1::> security ssh remove -vserver cluster1 -key-exchange-algorithms
diffie-hellman-group1-sha1,diffie-hellman-group-exchange-sha1 -ciphers
aes128-cbc,3des-cbc -mac-algorithms hmac-sha1-96,hmac-sha2-256
```

security ssh show

Display SSH configuration options

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

```
The ` security ssh show` command displays the configurations of the SSH
key exchange algorithms, ciphers, MAC algorithms, maximum authentication
retry count, host key algorithms and whether ``_ssh-rsa_`` signature
scheme is enabled for RSA keys in publickey algorithms, for the cluster
and Vservers. The SSH protocol uses a Diffie-Hellman based key exchange
method to establish a shared secret key during the SSH negotiation phrase.
The key exchange method specifies how one-time session keys are generated
for encryption and authentication and how the server authentication takes
place. ONTAP supports the ``_diffie-hellman-group-exchange-sha256_`` ,
``_diffie-hellman-group16-sha512_`` and ``_diffie-hellman-group18-
sha512_`` key exchange algorithms for SHA-2. ONTAP also supports the
``_diffie-hellman-group-exchange-sha1_`` , ``_diffie-hellman-group14-
sha1_`` , and ``_diffie-hellman-group1-sha1_`` key exchange algorithms for
SHA-1. ONTAP also supports ``_ecdh-sha2-nistp256_`` , ``_ecdh-sha2-
nistp384_`` , ``_ecdh-sha2-nistp521_`` and ``_curve25519-sha256_`` . ONTAP
also supports the AES and 3DES symmetric encryptions (also known as
ciphers) of the following types: ``_aes256-ctr_`` , ``_aes192-ctr_`` ,
``_aes128-ctr_`` , ``_aes256-cbc_`` , ``_aes192-cbc_`` , ``_aes128-cbc_``
, ``_aes128-gcm_`` , ``_aes256-gcm_`` and ``_3des-cbc_`` . ONTAP supports
MAC algorithms of the following types: ``_hmac-sha1_`` , ``_hmac-sha1-
96_`` , ``_hmac-md5_`` , ``_hmac-md5-96_`` , ``_umac-64_`` , ``_umac-
128_`` , ``_hmac-sha2-256_`` , ``_hmac-sha2-512_`` , ``_hmac-sha1-etm_`` ,
``_hmac-sha1-96-etm_`` , ``_hmac-sha2-256-etm_`` , ``_hmac-sha2-512-etm_``
, ``_hmac-md5-etm_`` , ``_hmac-md5-96-etm_`` , ``_umac-64-etm_`` and
``_umac-128-etm_`` . ONTAP supports host key algorithms of the following
types: ``_ecdsa-sha2-nistp256_`` , ``_ssh-rsa_`` and ``_ssh-ed25519_`` .
```

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.

[-vserver <Vserver Name>] - Vserver

Identifies the Vserver for which you want to display the SSH key exchange algorithms, ciphers, MAC algorithms, maximum authentication retry count, host key algorithms and whether *`ssh-rsa`* signature scheme is enabled for RSA keys in publickey algorithms.

[-key-exchange-algorithms <algorithm name>,…] - Key Exchange Algorithms

Displays the Vserver or Vservers that have the specified key exchange algorithms enabled.

[-ciphers <cipher name>,…] - Ciphers

Displays the Vserver or Vservers that have the specified ciphers enabled.

[-mac-algorithms <MAC name>,…] - MAC Algorithms

Displays the Vserver or Vservers that have the specified MAC algorithm or algorithms.

[-max-authentication-retry-count <integer>] - Max Authentication Retry Count

Displays Vservers with a matching maximum authentication retry count value. The default value of this parameter is *6* .

[-host-key-algorithms <HostKey Algorithms>,…] - Host Key Algorithms

Displays Vservers with matching host key algorithms.

[-is-rsa-in-publickey-algorithms-enabled {true|false}] - Is ssh-rsa in Publickey Algorithms Enabled

Identifies whether *ssh-rsa* signature scheme, which uses the SHA-1 hash algorithm, is enabled or disabled for RSA keys in publickey algorithms. The default value of this parameter is *true* .

Examples

The following command displays the enabled SSH ciphers, key exchange algorithms, MAC algorithms, host key algorithms, whether *ssh-rsa* signature scheme is enabled for RSA keys in publickey algorithms and maximum number of authentication retry count for the cluster and all Vservers. The cluster settings are used as the default for all newly created Vservers:

```
cluster1::> security ssh show

Is ssh-rsa   Max

in Publickey Auth
                            Key Exchange   MAC            Host Key
Algorithms   Retry
Vserver        Ciphers      Algorithms     Algorithms     Algorithms
Enabled      Count
-------------- ------------ -------------- -------------- --------------
------------ -----
cluster1       aes256-ctr,  curve25519-    hmac-sha1-etm, ecdsa-sha2-
false            4
               3des-cbc     sha256,        hmac-sha2-256- nistp256
                            diffie-        etm,
                            hellman-       hmac-sha2-512
                            group16-sha512
Is ssh-rsa   Max

in Publickey Auth
                            Key Exchange   MAC            Host Key
Algorithms   Retry
Vserver        Ciphers      Algorithms     Algorithms     Algorithms
Enabled      Count
-------------- ------------ -------------- -------------- --------------
------------ -----
vs1            aes256-ctr,  diffie-        hmac-sha2-256, ecdsa-sha2-
true             6
               aes192-ctr,  hellman-group- hmac-sha2-512, nistp256,
               aes128-ctr,  exchange-      hmac-sha2-256- ssh-rsa,
               aes128-gcm,  sha256,        etm,           ssh-ed25519
               aes256-gcm   ecdh-sha2-     hmac-sha2-512-
                            nistp256,      etm, umac-64,
                            ecdh-sha2-     umac-128,
                            nistp384,      umac-64-etm,
                            ecdh-sha2-     umac-128-etm
                            nistp521,
                            curve25519-
                            sha256
2 entries were displayed.
```

security ssl commands

security ssl modify

Modify the SSL configuration for HTTP servers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command modifies the configuration of encrypted HTTP (SSL) for Vservers in the cluster. Depending on the requirements of the individual node's or cluster's web services (displayed by the vserver services web show command), this encryption might or might not be used. If the Vserver does not have a certificate associated with it, SSL will not be available.

Parameters

`-vserver <Vserver Name>` - Vserver

Identifies a Vserver for hosting SSL-encrypted web services.

`[-ca <text>]` - Server Certificate Issuing CA

Identifies a Certificate Authority (CA) of a certificate to be associated with the instance of a given Vserver. If this parameter, along with serial, is omitted during modification, a self-signed SSL certificate can be optionally generated for that Vserver.

`[-serial <text>]` - Server Certificate Serial Number

Identifies a serial number of a certificate to be associated with the instance of a given Vserver. If this parameter, along with ca, is omitted during modification, a self-signed SSL certificate can be optionally generated for that Vserver.

`[-common-name <FQDN or Custom Common Name>]` - Server Certificate Common Name

Identifies the common name (CN) of a certificate to be associated with the instance of a given Vserver. This parameter becomes optional if serial and ca are specified. You can use the security certificate create and security certificate install commands to add new certificates to Vservers.

> The use of self-signed SSL certificates exposes users to man-in-the-middle security attacks. Where possible, obtain a certificate that is signed by a reputable certificate authority (CA) and use the security certificate install command to configure it before enabling SSL on a Vserver.

`[-server-enabled {true|false}]` - SSL Server Authentication Enabled

Defines the working condition of SSL server authentication in an instance of the Vserver. Any Vserver with a valid certificate of type server is server-enabled.

`[-client-enabled {true|false}]` - SSL Client Authentication Enabled

Defines the working condition of SSL client authentication in an instance of the Vserver. Any Vserver with a valid certificate of type client-ca is client-enabled. It can only be enabled if server-enabled is true.

`[-ocsp-enabled {true|false}]` - Online Certificate Status Protocol Validation Enabled

This parameter enables OCSP validation of the client certificate chain. When this parameter is enabled, certificates in the certificate chain of the client will be validated against an OCSP responder after normal verification (including CRL checks) has occurred. The OCSP responder used for validation process is either extracted from the certificate itself, or it is derived by configuration.

[-ocsp-default-responder <text>] - URI of the Default Responder for OCSP Validation

This parameter sets the default OCSP responder to use. If this parameter is not enabled, the URI given will be used only if no responder URI is specified in the certificate that are being verified.

[-ocsp-override-responder {true|false}] - Force the Use of the Default Responder URI for OCSP Validation

This parameter forces the configured default OCSP responder to be used during OCSP certificate validation, even if the certificate that is being validated references an OCSP responder.

[-ocsp-responder-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Timeout for OCSP Queries

Use this parameter to specify the timeout in seconds for OSCP responders. Specify zero for the minimum possible timeout. The default value is 10 seconds.

[-ocsp-max-response-age <integer_or_unlimited>] - Maximum Allowable Age for OCSP Responses (secs)

This parameter sets the maximum allowable age (freshness) in seconds for the OCSP responses. The default value for this parameter is unlimited, which does not enforce a maximum age and the OCSP responses are considered valid as long as their expiration date field is in the future.

[-ocsp-max-response-time-skew <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Maximum Allowable Time Skew for OCSP Response Validation

This parameter sets the maximum allowable time difference for OCSP responses (when validating their "ThisUpdate" and "NextUpdate" fields).

[-ocsp-use-request-nonce {true|false}] - Use a NONCE within OCSP Queries

This parameter determines whether the queries to the OCSP responders should contain a NONCE or not. By default, a query NONCE is always used and checked against the OCSP response. When the responder does not use NONCEs, this parameter should be disabled.

> A NONCE is a unique identifier included in each OCSP request or OCSP response to prevent a replay attack.

Examples

The following example enables SSL server authentication for a Vserver named vs0 with a certificate that has ca as www.example.com and serial as 4F4EB629.

```
cluster1::> security ssl modify -vserver vs0 -ca www.example.com -serial
4F4EB629 -server-enabled true
```

The following example disables SSL server authentication for a Vserver name vs0.

```
cluster1::> security ssl modify -vserver vs0 -server-enabled false
```

The following example enables SSL client authentication for a Vserver named vs0.

```
cluster1::> security ssl modify -vserver vs0 -client-enabled true
```

The following example disables SSL client authentication for a Vserver named vs0.

```
cluster1::> security ssl modify -vserver vs0 -client-enabled false
```

Related Links

- vserver services web show
- security certificate create
- security certificate install

security ssl show

Display the SSL configuration for HTTP servers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the configuration of encrypted HTTP (SSL) for Vservers in the cluster. Depending on the requirements of the individual node's or cluster's web services (displayed by the vserver services web show command), this encryption might or might not be used. If the Vserver does not have a certificate associated with it, SSL will not be available.

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.

| [-ocsp]

If you specify the `-ocsp` parameter, the command displays the Online Certificate Status Protocol configuration.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <Vserver Name>] - Vserver

Identifies a Vserver for hosting SSL-encrypted web services.

[-ca <text>] - Server Certificate Issuing CA

Filters the display of SSL configuration by specifying the Certificate Authority (CA) that issued the server certificate.

[-serial <text>] - Server Certificate Serial Number

Filters the display of SSL configuration by specifying the serial number of a server certificate.

[-common-name <FQDN or Custom Common Name>] - Server Certificate Common Name

Filters the display of SSL configuration by specifying the common name for the server certificate.

[-server-enabled {true|false}] - SSL Server Authentication Enabled

Filters the display of SSL configuration according to whether the SSL server authentication is enabled or disabled. Vservers have self-signed certificates automatically generated during their creation. These Vserver self-signed certificates are server-enabled by default.

[-client-enabled {true|false}] - SSL Client Authentication Enabled

Filters the display of SSL configuration according to whether the SSL client authentication is enabled or disabled. You can enable client authentication only when server authentication is enabled.

[-ocsp-enabled {true|false}] - Online Certificate Status Protocol Validation Enabled

Filters the display of SSL configuration when the Online Certificate Status Protocol validation is enabled.

[-ocsp-default-responder <text>] - URI of the Default Responder for OCSP Validation

Filters the display of SSL configuration according to the URI of the default responder for OCSP validation.

[-ocsp-override-responder {true|false}] - Force the Use of the Default Responder URI for OCSP Validation

Filters the display of SSL configuration, which forces the use of the default responder URI for OCSP validation.

[-ocsp-responder-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Timeout for OCSP Queries

Filters the display of SSL configuration according to the timeout for queries to OCSP responders.

[-ocsp-max-response-age <integer_or_unlimited>] - Maximum Allowable Age for OCSP Responses (secs)

Filters the display of SSL configuration according to the maximum allowable age (freshness) in seconds for the OCSP responses.

[-ocsp-max-response-time-skew <[<integer>d][<integer>h][<integer>m][<integer>s]>] - Maximum Allowable Time Skew for OCSP Response Validation

Filters the display of SSL configuration according to the maximum allowable time difference for OCSP responses (when validating their "ThisUpdate" and "NextUpdate" fields).

[-ocsp-use-request-nonce {true|false}] - Use a NONCE within OCSP Queries

Filters the display of SSL configuration by specifying whether the queries to the OCSP responders should contain a NONCE or not.

> A NONCE is a unique identifier included in each OCSP request or OCSP response to prevent a replay attack.

Examples

The following example displays the configured certificates for Vservers.

```
cluster1::security ssl> show
          Serial                                          Server  Client
Vserver   Number Common Name                              Enabled Enabled
--------- ------ ---------------------------------------- ------- -------
cluster1  516C3CB3
                 cluster1.company.com                     true    true
vs0       516816D4
                 vs0.company.com                          true    false
2 entries were displayed.
```

Related Links

- vserver services web show

security tpm commands

security tpm show

Display the status of TPM

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command displays information about the status of the Trusted Platform Module(TPM) device. By default, this command displays the following information:

- Node name
- Availability of the device
- State of the device, if available
- Firmware version
- Firmware upgrade counter

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>|local}] - Node

Selects the nodes that match this parameter value

[-is-available {yes|no}] - Is TPM Available?

Selects the nodes that match this parameter value.

- yes - The TPM device is mounted and available.
- no - The node does not support a TPM.

[-is-active {yes|no}] - Is TPM Active?

Selects the nodes that match this parameter value.

- yes - The TPM device is available and active.
- no - if `-is-available` parameter is *`yes`* , the TPM device is mounted and available but is not responding to TPM commands.

[-version <text>] - Firmware Version

Selects the nodes that match this firmware version.

[-upgrade-count <integer>] - Firmware Counter

Selects the nodes that match the given number of firmware upgrade tries left.

[-sym-key-size <integer>] - Size of Primary Symmetric Key

Selects the nodes that match the given symmetrice key size for the primary symmetric key.

Examples

```
cluster1::> security tpm show
                                          Firmware  Firmware
Node             Available?  Active?   Version    Counter
---------------- ----------  -------   --------  --------
node1            yes         yes       2.5             64
node2            yes         yes       2.5             64
2 entries were displayed.
```

security webauthn commands

security webauthn show

Show Webauthn Global Settings

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security webauthn show` displays information on the global FIDO2/WebAuthn settings.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that you specify.

| [-instance] }

Displays all the fields for the global WebAuthn settings.

[-vserver <vserver name>] - Vserver

Selects the WebAuthn settings that match this parameter value. If not specified, all cluster-level and Vserver-level settings are displayed.

[-timeout <integer>] - Timeout (ms)

Displays all the fields for the FIDO2/WebAuthn global settings.

[-attestation {none|indirect|direct|enterprise}] - Attestation Conveyance Type

Displays all the fields for the FIDO2/WebAuthn global settings.

[-require-rk {true|false}] - Require Resident Key

Displays all the fields for the FIDO2/WebAuthn global settings.

[-resident-key {required|preferred|discouraged}] - Resident Key

Displays all the fields for the FIDO2/WebAuthn global settings.

[-user-verification {required|preferred|discouraged}] - User Verification

Displays all the fields for the FIDO2/WebAuthn global settings.

[-scope {svm|cluster}] - Cluster or SVM

Displays all the fields for the FIDO2/WebAuthn global settings.

Examples

The example below displays information on FIDO2/WebAuthn global settings:

```
cluster1::> security webauthn show
Vserver: cluster1
               Timeout (ms): 600000
Attestation Conveyance Type: none
       Require Resident Key: false
               Resident Key: discouraged
          User Verification: discouraged
```

security webauthn credentials delete

Delete Webauthn Credentials Data

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security webauthn credentials delete` command deletes a previously registered FIDO2/WebAuthn credential. All users can delete their own registered FIDO2/WebAuthn credentials, regardless of their authorization role. Only users with *admin* privileges can delete FIDO2/WebAuthn credentials belonging to other users.

Parameters

-vserver <vserver name> - Vserver

This parameter optionally specifies the Vserver associated with the FIDO2/WebAuthn credential. If not specified, this defaults to the cluster Vserver.

-username <text> - Username

Name of the user whose credential is to be deleted.

-index <integer> - Index

The index of the FIDO2/WebAuthn credential to be deleted. This is used to identify the specific credential if the user has multiple FIDO2/WebAuthn credentials registered and only one of them is to be deleted.

-rpId <text> - Relying Party ID

The rpId of the FIDO2/WebAuthn credential to be deleted.

Examples

The following command, when executed by a user with admin privileges, deletes all registered FIDO2/WebAuthn credentials for the Administrative Vserver. When executed by users without admin privileges, the command only deletes that user's own registered credentials:

```
cluster1::> security webauthn credentials delete *
4 entries were deleted.
```

The following command, when executed by a user with admin privileges, deletes all registered FIDO2/WebAuthn credentials for the Administrative Vserver for the username *user1* :

```
cluster1::> security webauthn credentials delete -username user1 *
1 entry was deleted.
```

The following command, when executed by a user with admin privileges, deletes the registered FIDO2/WebAuthn credentials for the Administrative Vserver for the username *user1* at index *1* for the rpId *myontap.example.com* :

```
cluster1::> security webauthn credentials delete -username user1 -index 1
-rpId myontap.example.com
1 entry was deleted.
```

security webauthn credentials show

Show Webauthn Credentials Data

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security webauthn credentials show` displays information on the FIDO2/WebAuthn registered credentials.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that you specify.

| [-instance] }

Displays all the fields for the FIDO2/WebAuthn registered credentials.

[-vserver <vserver name>] - Vserver

Selects the FIDO2/WebAuthn registered credentials that match this parameter value. Only cluster-level credentials are supported at this time.

[-username <text>] - Username

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-index <integer>] - Index

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-rpId <text>] - Relying Party ID

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-credential-id-sha <text>] - SHA-256 Hashed Credential ID

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-publickey-algorithm <text>] - Public Key Algorithm

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-publickey <text>] - Public Key

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-display-name <text>] - Display Name

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-credential-type <public-key>] - Credential Type

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-rpName <text>] - Relying Party Name

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-creation-time <MM/DD/YYYY HH:MM:SS>] - Creation Time

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-last-used <MM/DD/YYYY HH:MM:SS>] - Last Used

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

[-scope {svm|cluster}] - Cluster or SVM

Selects the FIDO2/WebAuthn registered credentials that match this parameter value.

Examples

The example below displays information on the FIDO2/WebAuthn registered credentials:

```
cluster1::> security webauthn credentials show
Vserver: cluster1
                             Username: admin
                                Index: 0
                     Relying Party ID: myontap.example.com
          SHA-256 Hashed Credential ID: aaaaaaaaaaaaaaaaaaaaaaaaaaa
                  Public Key Algorithm: ES256
                            Public Key: abcd=
                          Display Name: admin
                       Credential Type: public-key
                    Relying Party Name: myontap.example.com
                         Creation Time: 6/10/2024 15:01:00
                             Last Used: 6/10/2024 15:01:00
```

security webauthn supported-algorithms show

Show WebAuthn Supported Algorithms

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `security webauthn supported-algorithms show` displays information on the FIDO2/WebAuthn supported algorithms.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that you specify.

| [-instance] }

Displays all the fields for the FIDO2/WebAuthn supported algorithms.

[-vserver <vserver name>] - Vserver

Selects the FIDO2/WebAuthn settings that match this parameter value. If not specified, all cluster-level and Vserver-level settings are displayed.

[-alg-name <text>] - Algorithm Name

Displays all the fields for the FIDO2/WebAuthn supported algorithms.

[-alg-id <integer>] - Algorithm ID

Displays all the fields for the FIDO2/WebAuthn supported algorithms.

[-alg-type <public-key>] - Algorithm Type

Displays all the fields for the FIDO2/WebAuthn supported algorithms.

[-scope {svm|cluster}] - Cluster or SVM

Displays all the fields for the FIDO2/WebAuthn supported algorithms.

Examples

The example below displays information on the FIDO2/WebAuthn supported algorithms:

```
cluster1::> security webauthn supported-algorithms show

Vserver: cluster1

ID         Name        Type
------     -------     -----
    -7     ES-256      public-key
   -35     ES-384      public-key
  -257     RS-256      public-key
3 entries were displayed.
```

snaplock commands

snaplock compliance-clock commands

snaplock compliance-clock initialize

Initializes the node ComplianceClock

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

` snaplock compliance-clock initialize` command is used to initialize System ComplianceClock from the system clock. User should ensure that the system clock is correct before initializing the System ComplianceClock. System ComplianceClock can be initialized multiple times as long as all nodes in the cluster are healthy, all volumes are in online state, no volumes are present in the volume recovery queue and there are no SnapLock volumes or volumes with "snapshot-locking-enabled" parameter set to true or S3 buckets with object locking enabled.

Parameters

`-node {<nodename>|local}` - Node

Specifies the name of the node on which System ComplianceClock needs to be initialized.

`[-force <true>]` - Forces Initialization

If you use this parameter, it will suppress the warning message during `snaplock compliance-clock initialize` operation.

Examples

```
cluster-1::> snaplock compliance-clock initialize -node node1

Warning: You are about to initialize the secure ComplianceClock of the
node
         node1 to the current value of the node's system clock. This
         procedure can be performed only once on a given node, so you
should
         ensure that the system time is set correctly before proceeding.
         The current node's system clock is: Wed Nov 26 16:18:30 IST 2014

Do you want to continue? {y|n}: y

cluster-1::>
```

snaplock compliance-clock show

Displays the node ComplianceClock

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The ` snaplock compliance-clock show ` command will display System ComplianeClock of the nodes in the cluster. It will display the following information:

- Node name
- ComplianceClock Time

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>|local}] - Node

If this parameter is specified, the command will display ComplianceClock for that particular node only.

[-time <text>] - ComplianceClock Time of the Node

If this parameter is specified, the command will display nodes having the same `-time` value.

Examples

```
cluster1::> snaplock compliance-clock show
        Node                ComplianceClock Time
        ------------------- ------------------------------------
        node1               Mon Jan 12 11:34:15 IST 2015 +05:30
        node2               Mon Jan 12 11:34:10 IST 2015 +05:30
        2 entries were displayed.
```

```
cluster1::> snaplock compliance-clock show -node node1
        Node                ComplianceClock Time
        ------------------- ------------------------------------
        node1               Mon Jan 12 11:34:45 IST 2015 +05:30
```

snaplock compliance-clock ntp modify

Modify SnapLock ComplianceClock synchronization setting

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `snaplock compliance-clock ntp modify` command modifies the option to enable or disable the SnapLock ComplianceClock synchronization with the system time. The ComplianceClock is synchronized only when an NTP server has been configured so that the system time follows the NTP time and the skew between the ComplianceClock time and the system time is greater than 1 day.

Parameters

`[-is-sync-enabled {true|false}]` - Enable ComplianceClock sync to NTP system time (privilege: advanced)

Specifies whether synchronization should be enabled or not. This is a cluster wide option.

Examples

```
cluster1::> snaplock compliance-clock ntp modify -is-sync-enabled true
```

snaplock compliance-clock ntp show

Display SnapLock ComplianceClock synchronization setting

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The ` snaplock compliance-clock ntp show ` command will display ComplianceClock synchronization setting. It will display the following information:

- is-sync-enabled - Displays if the option to synchronize the ComplianceClock with system time has been enabled or not.

Examples

```
cluster1::> snaplock compliance-clock ntp show
Enable clock sync to NTP system time: true
```

snaplock event-retention commands

snaplock event-retention abort

Abort an Event Based Retention policy operation.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention abort` is used to abort an ongoing Event Based Retention (EBR) operation. This command only aborts the operations that have not yet completed. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the vserver on which the EBR operation is running.

`-operation-id <integer>` - Operation ID

Specifies the operation ID of the EBR operation that needs to be aborted.

Examples

The following example aborts an ongoing EBR operation with operation-id *`16842754`* :

```
vs1::> snaplock event-retention abort -operation-id 16842754
            vs1::>
```

snaplock event-retention apply

Apply an Event Based Retention policy on all files within a user specified path.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention apply` command starts a new operation to apply the specified Event Based Retention (EBR) policy to all files in the specified path. If a file is a regular file, it will be made a WORM file and retained for a retention-period as defined by the specified policy name. If a file is already WORM, its retention time will be extended to a retention-period as defined by the specified policy name, starting from the current time. The retention time of a file will be extended only if the file's current retention time is less than the new retention time value to be set. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver which has the EBR policy defined to be applied on one or more files.

`-policy-name <text>` - Policy Name

Specifies the name of the EBR policy to be applied on one or more files.

`-volume <volume name>` - Volume

Specifies the name of the SnapLock volume containing a file path or a directory path as specified by the path parameter. The specified EBR policy is applied to one or more files depending on the value of path.

`-path <text>` - Path

Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory. The EBR policy is applied to all files under the specified path. To apply the EBR policy to all files in a volume, specify the path as "/".

Examples

The following example starts an EBR operation to apply a policy on files for specified volume:

```
vs1::> snaplock event-retention apply -policy-name p1 -volume slc -path /
                SnapLock event based retention operation is queued. Run
"snaplock event-retention show -operation-id 16842754 -instance" to view
the operation status.
```

snaplock event-retention show-vservers

Show Vservers with SnapLock Event Based Retention policies

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention show-vservers` command is used to display the Vservers that have SnapLock Event Based Retention (EBR) policies created.

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.

Examples

The following example displays all Vservers that have SnapLock EBR policies:

```
cluster-1::*> snaplock event-retention show-vservers
        Vserver
        -----------------
        vs1
```

snaplock event-retention show

Show status of Event Based Retention operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention show` command displays the status of an Event Based Retention (EBR) operation. Information about completed operations will be cleaned up after an hour after completion. Only a

user with security login role *vsadmin-snaplock* is allowed to perform this operation.

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.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, the command displays all EBR operations that match the specified Vserver.

[-operation-id <integer>] - Operation ID

If this parameter is specified, the command displays all EBR operations that match the specified operation ID.

[-volume <volume name>] - Volume Name

If this parameter is specified, the command displays all EBR operations that match the specified volume. The parameter specifies the volume on which EBR operation is running or has completed.

[-path <text>] - Path

If this parameter is specified, the command displays all EBR operations that match the specified path. The parameter specifies the path on which EBR operation is running or has completed.

[-policy-name <text>] - Policy Name

If this parameter is specified, the command displays all EBR operations that match the specified policy name. The parameter specifies the EBR policy name.

[-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}] - Retention Period

If this parameter is specified, the command displays all EBR operations that match the specified retention period. The parameter specifies the retention period of the EBR policy.

[-num-files-processed <integer>] - Number of Files Processed

If this parameter is specified, the command displays all EBR operations that match the specified number of processed files. The parameter specifies the number of files on which EBR policy was applied successfully.

[-num-files-failed <integer>] - Number of Files Failed

If this parameter is specified, the command displays all EBR operations that match the specified number of failed files. The parameter specifies the number of files on which the application of EBR policy failed.

[-num-files-skipped <integer>] - Number of Files Skipped

If this parameter is specified, the command displays all EBR operations that match the specified number of skipped files. The parameter specifies the number of files on which the application of EBR policy was skipped. A file that is under legal-hold will be skipped. If the retention time of a file is being shortened as a result of applying the EBR policy, that file will also be skipped.

[-num-inodes-ignored <integer>] - Number of Inodes Ignored

If this parameter is specified, the command displays all EBR operations that match the specified number of ignored inodes. The parameter specifies the number of inodes on which the application of EBR policy was not attempted because they were not regular files.

[-operation-status {Unknown|In-Progress|Failed|Aborting|Completed}] - Operation Status

If this parameter is specified, the command displays all EBR operations that match the specified operation status. The parameter specifies the operation status of an EBR operation.

[-status-details <text>] - Status Details

If this parameter is specified, the command displays all EBR operations that match the specified status details. The parameter specifies the status details of an EBR operation.

Examples

The following examples show the status of EBR operations for Vserver "vs1" and volume "slc" and the status of event-retention operation for operation ID *16842753* respectively.

```
vs1::*> snaplock event-retention operation show -volume slc
            Operation ID   Vserver         Volume          Operation
Status
            -------------- --------------- ---------------
----------------
            16842753       vs1             slc             Completed
            16842754       vs1             slc             In-progress
vs1::*> snaplock event-retention operation show -operation-id 16842753
Operation ID: 16842753
                    Vserver: vs1
                     Volume: slc
                       Path: /vol/slc/d1
                Policy Name: p1
           Retention Period: 10 years
  Number of Files Processed: 50
     Number of Files Failed: 0
   Number of Inodes Ignored: 2
           Operation Status: Completed
             Status Details: No error
```

snaplock event-retention policy create

Create SnapLock Event Based Retention policies for a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention policy create` command is used to create Event Based Retention (EBR) policies for a Vserver. A policy consists of a *`policy-name`* and a *`retention-period`* . Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver for which a policy needs to be created.

`-name <text>` - Policy Name

Specifies the name of the EBR policy to be created.

`-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}` - Event Retention Period

Specifies the retention period for an EBR policy.

Examples

The following example creates a new EBR policy "p1" for Vserver "vs1" with a retention period of *`"10 years"`* :

```
vs1::> snaplock event-retention policy create -name p1 -retention-period
"10 years"
```

snaplock event-retention policy delete

Delete SnapLock Event Based Retention policies for a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention policy delete` command is used to delete Event Based Retention (EBR) policies for a Vserver. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

If this parameter is specified, the command deletes all EBR policies that match the specified Vserver.

`-name <text>` - Policy Name

If this parameter is specified, the command deletes all EBR policies that match the specified *`name`* .

Examples

The following example deletes retention policy "p1" for Vserver "vs1":

```
vs1::> snaplock event-retention policy delete -name p1
```

snaplock event-retention policy modify

Modify SnapLock Event Based Retention policies for a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention policy modify` command is used to modify the retention period of an Event Based Retention (EBR) policy for a Vserver. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver for which retention period of a policy needs to be modified.

`-name <text>` - Policy Name

Specifies the name of the EBR policy for which the retention period needs to be modified.

`[-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}]` - Event Retention Period

Specifies the new value of retention period.

Examples

The following example modifies the retention period of policy "p1" for Vserver "vs1" to *`"5 years"`* :

```
vs1::> snaplock event-retention policy modify -name p1 -retention-period
"5 years"
```

snaplock event-retention policy show

Show SnapLock Event Based Retention policies for a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock event-retention policy show` command is used to show Event Based Retention (EBR) policies for a Vserver. A policy consists of a *`policy-name`* and a *`retention-period`* . The command output depends on the parameter or parameters specified. If no parameters are specified, all policies for all vservers will be displayed. If one or more parameters are specified, only those entries matching the specified values will be displayed. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

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.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, the command displays all EBR policies that match the specified Vserver.

[-name <text>] - Policy Name

If this parameter is specified, the command displays all EBR policies that match the specified *`name`* .

[-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}] - Event Retention Period

If this parameter is specified, the command displays all EBR policies that match the specified *`retention-period`*.

Examples

The following example displays all event-retention policies for vserver "vs1":

```
vs1::> snaplock event-retention policy show
            Vserver           Name                Retention Period
            ----------------- ------------------- --------------------
            vs1               p1                  10 years
            vs1               p2                  5 years
```

snaplock legal-hold commands

snaplock legal-hold abort

Abort Snaplock legal-hold operation.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock legal-hold abort` is used to abort an ongoing legal-hold operation. The type of legal-hold operations that can be aborted using this command are begin, end and dump-files. This command only aborts operations that have not yet completed. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the vserver on which the legal-hold operation is running.

`-operation-id <integer>` - Operation ID

Specifies the operation ID of the legal-hold operation to be aborted.

Examples

The following example aborts an ongoing legal-hold operation with operation-id *`16842754`* :

```
vs1::> snaplock legal-hold abort -operation-id 16842754
            vs1::>
```

snaplock legal-hold begin

Starts an operation to place files under legal-hold in the user specified path on a SnapLock compliance volume.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock legal-hold begin` command is used to place specified file or files under legal-hold for a given litigation. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver which owns the volume. The specified file or files to be placed under legal-hold reside on this volume.

`-litigation-name <text>` - Litigation Name

Specifies the name of the litigation for which the file or files have to be placed under legal-hold.

`-volume <volume name>` - Volume

Specifies the name of the SnapLock compliance volume on which the file or files to be placed under legal-hold reside.

`-path <text>` - Path

Specifies a path relative to the volume root. The path can be either a file path of the single file to be placed under legal-hold or a directory path where all regular files under it must be placed under legal-hold.

Examples

The following example starts a legal-hold begin operation on file *`file1`* in volume *`slc_vol1`* :

```
vs1::> snaplock legal-hold begin -litigation-name litigation1 -volume
slc_vol1 -path /file1
SnapLock legal-hold begin operation is queued. Run "snaplock legal-hold
show -operation-id 16842773 -instance" to view the operation status.
```

The following example starts a legal-hold begin operation on all files in the volume *slc_vol1* :

```
vs1::> snaplock legal-hold begin -litigation-name litigation1 -volume
slc_vol1 -path /
SnapLock legal-hold begin operation is queued. Run "snaplock legal-hold
show -operation-id 16842775 -instance" to view the operation status.
```

snaplock legal-hold dump-files

Dump list of files under legal-hold to specified output path.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock legal-hold dump-files` is used to dump the list of files under legal-hold for a given vserver, volume and litigation to an auto-generated file in the user specified path. Only a user with security login role *vsadmin-snaplock* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver for which the list of files under legal-hold is to be dumped.

`-litigation-name <text>` - Litigation Name

Specifies the name of the litigation for which the list of files under legal-hold is to be dumped.

`-volume <volume name>` - Volume Name

Specifies the name of the SnapLock compliance volume for which the list of files under legal-hold is to be dumped.

`-output-volume <volume name>` - Output Volume Name

Specifies the name of the output volume containing the output directory path where the list of files under legal-hold is to be dumped. The output volume must be a regular read-write volume.

`-output-directory-path <text>` - Path Relative to Output Volume Root

Specifies the output directory path relative to the output volume root, where the list of files under legal-hold is to be dumped. The output directory path should be of the form "/directory-path". If output needs to be dumped on the volume root, specify the path as "/".

Examples

The following example starts a legal-hold dump-files operation:

```
vs1::> snaplock legal-hold dump-files -volume vol1_slc -litigation-name
lit1 -output-volume vol1 -output-directory-path /d1
            SnapLock legal-hold dump-files operation is queued. Run
"snaplock legal-hold show -operation-id 16842754 -instance" to view the
operation status.
            vs1::>
```

snaplock legal-hold dump-litigations

Dump list of litigations for a given Vserver to specified output path.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock legal-hold dump-litigations` is used to dump the list of litigations for a given vserver to a user specified path. Under the user specified path, we create a directory with a unique name. Under the user specified path, a directory with an auto-generated name is created. Under this directory, multiple files are created. Each file represents a unique litigation name that was found in the given vserver. Each file contains a list of volume names that have files under legal-hold for that given litigation. For example, if the file name is "lit1" and the contents of the file are "volume1" and "volume2", then it indicates that both these volumes have files under legal-hold for litigation "lit1". Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver for which the list of litigations is to be dumped.

`[-volume <volume name>]` - Volume Name

If this parameter is specified, the command displays the list of litigations for volume that matches the specified value. The volume must be of type SnapLock compliance.

`-output-volume <volume name>` - Output Volume Vame

Specifies the name of the output volume containing the output directory path where the list of litigations is to be dumped. The output volume must be a regular read-write volume.

`-output-directory-path <text>` - Path Relative to Output Volume Root

Specifies the output directory path relative to the volume root, where the list of litigations is to be dumped. The output directory path should be of the form "/directory-path". If output needs to be dumped to the volume root, specify the path as "/".

Examples

The following example starts a legal-hold dump-litigations job:

```
vs1::> snaplock legal-hold dump-litigations -output-volume vol1 -output
-directory-path /d1
            Dump Litigations job for Vserver "vs1" has been queued. Run
"job show -id 22 -instance" to view the status.
            vs1::>
```

snaplock legal-hold end

Starts an operation to release legal-hold on files in the user specified path on a SnapLock compliance volume.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock legal-hold end` command is used to release legal-hold on specified file or files for a given litigation. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver which owns the volume. The specified file or files to be released from legal-hold reside on this volume.

`-litigation-name <text>` - Litigation Name

Specifies the name of the litigation for which the file or files have to release from legal-hold.

`-volume <volume name>` - Volume

Specifies the name of the SnapLock compliance volume on which the file or files to be released from legal-hold reside.

`-path <text>` - Path

Specifies a path relative to the volume root. The path can be either a file path of the single file to be released from legal-hold or a directory path where all regular files under it must be released from legal-hold.

Examples

The following example starts a legal-hold end operation on file *`file1`* in volume *`slc_vol1`* :

```
vs1::> snaplock legal-hold end -litigation-name litigation1 -volume
slc_vol1 -path /file1
SnapLock legal-hold end operation is queued. Run "snaplock legal-hold show
-operation-id 16842773 -instance" to view the operation status.
```

The following example starts a legal-hold end operation on all files in the volume *`slc_vol1`* :

```
vs1::> snaplock legal-hold end -litigation-name litigation1 -volume
slc_vol1 -path /
SnapLock legal-hold end operation is queued. Run "snaplock legal-hold show
-operation-id 16842775 -instance" to view the operation status.
```

snaplock legal-hold show

Show status of a legal-hold operation.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock legal-hold show` command displays the status of a legal-hold operation. Information about completed operations will be cleaned up after an hour of completion. Only a user with security login role *`vsadmin-snaplock`* is allowed to perform this operation.

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.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, the command displays all legal-hold operations that match the specified Vserver.

[-operation-id <integer>] - Operation ID

If this parameter is specified, the command displays all legal-hold operations that match the specified operation ID.

[-volume <volume name>] - Volume Name

If this parameter is specified, the command displays all legal-hold operations that match the specified volume. The parameter specifies the volume on which legal-hold operation is running or has completed.

[-path <text>] - Path

If this parameter is specified, the command displays all legal-hold operations that match the specified path. The parameter specifies the path on which legal-hold operation is running or has completed.

[-litigation-name <text>] - Litigation Name

If this parameter is specified, the command displays all legal-hold operations that match the specified litigation name. The parameter specifies the legal-hold litigation name.

[-operation-type {unknown|begin|end|dump-files}] - Operation Type

If this parameter is specified, the command displays all legal-hold operations that match the specified

operation type. The parameter specifies the type of legal-hold operation.

`[-operation-status {Unknown|In-Progress|Failed|Aborting|Completed}]` - Operation Status

If this parameter is specified, the command displays all legal-hold operations that match the specified operation status. The parameter specifies the status of legal-hold operation.

`[-num-files-processed <integer>]` - Number of Files Processed

If this parameter is specified, the command displays all legal-hold operations that match the specified number of files processed. The parameter specifies the number of files on which legal-hold operation was successful.

`[-num-files-failed <integer>]` - Number of Files Failed

If this parameter is specified, the command displays all legal-hold operations that match the specified number of files failed. The parameter specifies the number of files on which legal-hold operation failed.

`[-num-files-skipped <integer>]` - Number of Files Skipped

If this parameter is specified, the command displays all legal-hold operations that match the specified number of files skipped.

The parameter specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation or if it is a hard link to a file that is already under hold for a given litigation.

`[-num-inodes-ignored <integer>]` - Number of Inodes Ignored

If this parameter is specified, the command displays all legal-hold operations that match the specified number of inodes ignored. The parameter specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.

`[-status-details <text>]` - Status Details

If this parameter is specified, the command displays all legal-hold operations that match the specified status details. The parameter specifies the status details of an legal-hold operation.

Examples

The following examples show the status of legal-hold operations for Vserver *`vs1`* and volume *`slc_vol1`* and the status of legal-hold operation for operation ID *`16842786`* respectively:

```
vs1::> snaplock legal-hold show -volume slc_vol1
                                                Operation
Operation     Operation ID    Vserver   Volume  Status
------------- --------------- --------- ------- ---------
begin         16842784        vs1       slc_vol1
                                                Completed
begin         16842786        vs1       slc_vol1
                                                Completed
begin         16842788        vs1       slc_vol1
                                                In-Progress
dump-files    16842790        vs1       slc_vol1
                                                Completed
end           16842794        vs1       slc_vol1
                                                Completed
5 entries were displayed.

vs1::> snaplock legal-hold show -operation-id 16842786
Vserver: vs1
                      Volume: slc_vol1
                Operation ID: 16842786
             Litigation Name: litigation1
                        Path: /
              Operation Type: begin
                      Status: Completed
   Number of Files Processed: 100
      Number of Files Failed: 15
     Number of Files Skipped: 20
    Number of Inodes Ignored: 0
              Status Details: No error
```

snaplock log commands

snaplock log create

Create audit log configuration for a Vserver.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

```
The `snaplock log create` command is used to create a SnapLock log
configuration for the Vserver. A SnapLock log configuration consists of
volume to store the log, the maximum size of the log file, and the default
period of time for which the log file should be retained.
```

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver for which the configuration needs to be created.

`-volume <volume name>` - Log Volume Name

Specifies the name of the volume that is used for logging. This must be a SnapLock Compliance volume.

`[-max-log-size {<integer>[KB|MB|GB|TB|PB]}]` - Maximum Size of Log File

Specifies the maximum size of the log file. Once a log file reaches this limit, it is archived and a new log file is created. This parameter is optional. The default value is *`10MB`* .

`[-default-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}]` - Default Log Record Retention Period

Specifies the default period of time a record (which is logged) is retained. This parameter is optional. The default value is *`"6 months"`* .

Examples

```
cluster1::> snaplock log create -volume vol1 -max-log-size 50MB -default
-retention-period "1 year" -vserver vs1
        [Job 47] Job succeeded: SnapLock log created for Vserver "vs1".
```

snaplock log delete

Delete audit log configuration for a Vserver.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock log delete` command deletes the SnapLock log configuration associated with the Vserver. This command closes all the active log files in the log volume and mark the volume as disabled for SnapLock logging.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver whose SnapLock log configuration is deleted.

Examples

```
cluster1::> snaplock log delete -vserver vs1
        [Job 47] Job succeeded: SnapLock log deleted for Vserver "vs1".
```

snaplock log modify

Modify audit log configuration for a Vserver.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock log modify` command modifies the SnapLock log configuration of the Vserver. Log volume, maximum size of log file, and default retention period can be modified. If the log volume is modified, then the active log files in the existing log volume is closed and the log volume is marked as disabled for logging. The new log volume is enabled for logging.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver for which the SnapLock log configuration needs to be modified.

`[-volume <volume name>]` - Log Volume Name

Specifies the new log volume that is configured for this Vserver for logging.

`[-max-log-size {<integer>[KB|MB|GB|TB|PB]}]` - Maximum Size of Log File

Specifies the new value for maximum log file size.

`[-default-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}]` - Default Log Record Retention Period

Specifies the new value for default retention period.

Examples

```
cluster1::> snaplock log modify -volume vol1 -vserver vs1 -max-log-size
15MB
        [Job 48] Job succeeded: SnapLock log modified for Vserver "vs1".
```

snaplock log show

Display audit log configuration.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock log show` command displays the following information about the SnapLock log infrastructure:

- Vserver name
- Volume name
- Maximum log size
- Default retention period

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.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, the command displays the log information for Vservers that match the specified value.

[-volume <volume name>] - Log Volume Name

If this parameter is specified, the command displays the log configuration for volumes that match the specified value.

[-max-log-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Size of Log File

If this parameter is specified, the command displays the log configuration with a matching `-max-log-size` value.

[-default-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}] - Default Log Record Retention Period

If this parameter is specified, the command displays the log configuration with a matching `-default-retention-period` value.

Examples

```
cluster1::> snaplock log show -vserver vs1
Vserver Name                        : vs1
      Log Volume Name                     : 15MB
      Maximum Size of Log File            : 15MB
      Default Log Record Retention Period : 6 months
```

```
cluster1::> snaplock log show
        Vserver           Volume              Maximum Size
Retention Period
        ----------------- ------------------- -------------------
--------------------
        vs1               vol1                15MB                6 months
```

snaplock log file archive

Archive Active Log Files in Log Volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock log file archive` command archives the currently active log file by closing it and creating a new active log file. If base-name is not provided, the command archives all active log files associated with the Vserver. Otherwise, the command archives the active log file associated with the base-name provided.

Parameters

`-vserver <vserver name>` - Vserver Name

Specifies the name of the Vserver for which active log files need to be archived.

`[-base-name {privileged-delete | system | legal-hold}]` - Base Name of Log File

Specifies the log base-name, whose active log file needs to be archived. The base-name is the name of the source of log records. Valid base-names are *`system`*, *`privileged-delete`* and *`legal-hold`*. Each base-name has its own directory in which log files containing log records generated by base-name are stored.

Examples

```
cluster1::> snaplock log archive -vserver vs1
        [Job 48] Job succeeded: SnapLock log archived for Vserver "vs1".
```

snaplock log file show

Display audit log file information.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snaplock log file show` command displays the following information about the log files:

- Vserver name
- Volume name
- File path
- Expiry time of the file
- File size

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.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, then log files in the Vserver that match the specified value is displayed.

[-base-name {privileged-delete | system | legal-hold}] - Base Name of Log File

If this parameter is specified, then the log files having a matching `-base-name` is displayed.

[-volume <volume name>] - Log Volume Name

If this parameter is specified, then the log files in volumes that match the specified value are shown.

[-file-path <text>] - Log File Path

If this parameter is specified, then the log files that match the specified value are displayed.

[-expiry-time <text>] - Log File Expiry Time

If this parameter is specified, then the log files having a matching `-expiry-time` value are displayed.

[-file-size {<integer>[KB|MB|GB|TB|PB]}] - File Size

If this parameter is specified, then the log files having a matching `-file-size` value are displayed.

Examples

```
cluster1::> snaplock log file show
        Vserver           Volume               Base Name           File
Path
        ----------------- -------------------- --------------------
--------------------
        vs1               vol1                 system
/vol/vol1/snaplock_log/system_logs/20160120_183756_GMT-present
```

```
cluster1::> snaplock log file show -vserver vs1 -base-name system
Vserver     : vs1
        Volume      : vol1
        Base Name   : system
        File Path   :
/vol/vol1/snaplock_log/system_logs/20160120_183756_GMT-present
        Expiry Time : Wed Jul 20 18:37:56 GMT 2016
        File Size   : 560B
```

snapmirror commands

snapmirror abort

Abort an active transfer

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror abort` command stops SnapMirror transfers that might have started and not completed. A SnapMirror transfer is an operation on a given SnapMirror relationship and the relationship is identified by its destination endpoint, which can be a volume, a Vserver, an Application Consistency Group or a non-Data ONTAP endpoint. You identify the SnapMirror relationship with this command and the command aborts the transfer for the relationship. For load-sharing mirrors, the command also aborts transfers for other relationships that are part of the same load-sharing set. For SolidFire destination endpoints, the `snapmirror abort` command is only supported if the endpoint is in a SnapMirror relationship.

Load-sharing mirrors are either up to date and serving data to clients, or they are lagging and not serving data to clients. If the `snapmirror abort` command identifies an up-to-date load-sharing mirror, then SnapMirror transfers to the up-to-date load-sharing mirror and associated up-to-date load-sharing mirrors in the set of load-sharing mirrors are aborted. If the `snapmirror abort` command identifies a lagging load-sharing mirror, then only the SnapMirror transfer associated with the lagging load-sharing mirror is aborted.

After the `snapmirror abort` command successfully completes its operation, the volume on the receiving side of the transfer might contain a restart checkpoint. The restart checkpoint can be used by a subsequent transfer to restart and continue the aborted SnapMirror transfer.

This command is supported for SnapMirror relationships with the field *`"Relationship Capability"`* showing as either *`"8.2 and above"`* or *`"Pre 8.2"`* in the output of the snapmirror show command.

This command is not supported for SnapMirror active sync relationships with policy of type *`automated-failover`* or *`automated-failover-duplex`* in admin privilege level.

The use of wildcards in parameter values is not supported from the source Vserver or cluster for relationships with *`"Relationship Capability"`* of *`"8.2 and above"`* .

You can use this command from the source or the destination Vserver or cluster for FlexVol volume relationships.

For SnapMirror Synchronous relationships, this command aborts any ongoing transfer and takes the relationship OutOfSync. This can result in primary client IO failure for relationships with a policy of type *`strict-sync-mirror`* . Instead, the best practice recommendation is to use the snapmirror quiesce command.

Beginning with ONTAP 9.12.1, for SnapMirror Synchronous relationships on platforms that support NDO, `snapmirror abort` triggered in *`"InSync"`* state will not change the relationship status to *`"OutOfSync"`* . Instead, it will internally trigger the feature *`"Fast Resync"`* which will try to keep the relationship status *`"InSync"`* .

For Vserver SnapMirror relationships, this command must be run only from the cluster containing the destination Vserver.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name*.

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name*.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

`-destination-vserver <vserver name>` - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

`-destination-volume <volume name>` - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

`[-h, -hard <true>]` - Discard Restart Checkpoint

If this option is specified true, the restart checkpoint is discarded and the destination volume is restored to the last snapshot that was successfully transferred. You can use the `-hard` option to discard the restart checkpoint of a previous transfer attempt which forces the subsequent transfer to start with a fresh snapshot on the destination volume. This option can only be used from the destination Vserver or cluster. This parameter is not supported for relationships with non-Data ONTAP endpoints.

`[-w, -foreground <true>]` - Foreground Process

This specifies whether the operation runs as a foreground process. If this parameter is specified, the default setting is *true* (the operation runs in the foreground). When set to *true*, the command will not return until the process completes. This parameter is only applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"*.

Examples

To stop the active SnapMirror replication to the destination volume *vs2.example.com:dept_eng_dp_mirror1*, type the following command:

```
vs2.example.com::> snapmirror abort -destination-path
       vs2.example.com:dept_eng_dp_mirror1
```

For relationships with *"Relationship Capability"* of *"Pre 8.2"*, to stop the active SnapMirror replication to the destination volume *cluster2://vs2.example.com/dept_eng_dp_mirror1*, type the following command:

```
cluster2::> snapmirror abort -destination-path
         cluster2://vs2.example.com/dept_eng_dp_mirror1
```

To stop the active SnapMirror replication to the destination Vserver *dvs1.example.com*, type the following command:

```
cluster2::> snapmirror abort -destination-path
         dvs1.example.com:
```

To stop the active SnapMirror replication to the destination Application Consistency Group *cg_dst* in Vserver *vs2.example.com* , type the following command:

```
cluster2::> snapmirror abort -destination-path
         vs2.example.com:/cg/cg_dst
```

Related Links

- snapmirror show
- snapmirror quiesce

snapmirror break

Make SnapMirror destination writable

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror break` command breaks a SnapMirror relationship between a source and destination endpoint of a data protection mirror. The destination endpoint can be a Vserver, volume, Application Consistency Group or SolidFire endpoint. When ONTAP breaks the relationship, if the endpoint is a volume, Application Consistency Group or SolidFire endpoint, the destination is made read/write and can diverge from the source volume, client redirection is turned off on the destination, the restart checkpoint is cleared, and the clients can see the latest snapshot. If the endpoint is a Vserver, the subtype of the destination Vserver is changed to *default* , volumes in the destination Vserver are made read/write and the clients can now access the Vserver namespace for modifications. For SolidFire destination endpoints, the `snapmirror break` command is only supported if the endpoint is in a SnapMirror relationship. For an Application Consistency Group the item volumes are changed to read/write volumes and can diverge from the corresponding items on the source.

Subsequent manual or scheduled SnapMirror updates to the broken relationship will fail until the SnapMirror relationship is reestablished using the snapmirror resync command.

This command applies to data protection mirrors. For vault relationships, this command is only intended for use when preparing for a ONTAP revert operation (see the `-delete-snapshots` parameter in advanced privilege level). This command is not intended for use with load-sharing mirrors.

For relationships with a policy of type *strict-sync-mirror* or *sync-mirror* , the relationship must be *Quiesced* before running the `snapmirror break` command.

This command is supported for SnapMirror relationships with the field *"Relationship Capability"* showing as either *"8.2 and above"* or *"Pre 8.2"* in the output of the snapmirror show command.

The `snapmirror break` command must be used from the destination Vserver or cluster.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name*.

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name*.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

`-destination-vserver <vserver name>` - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

`-destination-volume <volume name>` - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

`[-f, -force <true>]` - Force

If this parameter is specified, the command proceeds without prompting for confirmation.

`[-w, -foreground <true>]` - Foreground Process

This specifies whether the operation runs as a foreground process. If this parameter is specified, the default setting is *true* (the operation runs in the foreground). When set to *true* , the command will not return until the process completes. This parameter is only applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"*.

`[-delete-snapshots <true>]` - Delete Snapshots for Revert (privilege: advanced)

Using this parameter causes break to delete snapshots on a vault destination so that the system can be reverted. Note that the only snapshots that will be deleted are those that were created with the current version of ONTAP. Any snapshots that might be present created with a different version will not be deleted. This parameter is not supported for SnapLock Compliance destinations.

`[-s, -restore-destination-to-snapshot <text>]` - Restore Destination to Snapshot Copy

This optional parameter specifies the snapshot to which the destination volume is restored after a successful break operation. If the parameter is not specified, the destination is restored to the latest snapshot. This parameter is not supported for Vserver, Application Consistency Group or FlexGroup relationships or SnapLock Compliance destinations.

`[-recover <true>]` - Recover (privilege: advanced)

When a SnapMirror break operation fails on a FlexGroup or Application Consistency Group relationship, a subset of the destination FlexGroup or Application Consistency Group constituents could have been made writable and subsequently user data could have been written to these constituents. To recover from this failure, you can execute the `snapmirror break` command again specifying the `-recover` parameter. All constituents will be restored to the latest snapshot, and any writes to the read-write constituents will be lost. This parameter is applicable only for SnapMirror relationships with FlexGroup and or Application Consistency Group endpoints.

Examples

To stop the SnapMirror replication to the destination volume *vs2.example.com:dept_eng_dp_mirror1* , type the following command:

```
vs2.example.com::> snapmirror break -destination-path
vs2.example.com:dept_eng_dp_mirror1
```

For relationships with *"Relationship Capability"* of *"Pre 8.2"* , to stop the SnapMirror replication to the destination volume *cluster2://vs2.example.com/dept_eng_dp_mirror1* , type the following command:

```
cluster2::> snapmirror break
         -destination-path cluster2://vs2.example.com/dept_eng_dp_mirror1
```

To stop replication to the destination Vserver *dvs1.example.com* of a Vserver SnapMirror relationship, type the following command:

```
cluster2::> snapmirror break -destination-path dvs1.example.com:
```

To stop SnapMirror replication to the destination Application Consistency Group *app_cg_dst* in Vserver *vs2.example.com* , type the following command:

```
vs2.example.com::> snapmirror break -destination-path
       vs2.example.com:/cg/app_cg_dst
```

Related Links

- snapmirror resync
- snapmirror show

snapmirror create

Create a new SnapMirror relationship

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror create` command creates a SnapMirror relationship between a source and destination endpoint. You can use this command to create a data protection relationship, an extended data protection relationship, or a load-sharing relationship between FlexVol volumes. You can also use it to create a data protection relationship between Vservers. A SnapMirror relationship between Vservers can only be created if the system containing the source Vserver is also running ONTAP 8.3 or later. You can also use the `snapmirror create` command to create an extended data protection relationship between FlexGroup volumes. FlexGroup volumes only support extended data protection relationships. A SnapMirror relationship between FlexGroup volumes is only supported if the system containing the source FlexGroup volume is also running ONTAP 9.1.0 or later. The source or destination of a FlexGroup SnapMirror relationship cannot be the source or destination of any non-FlexGroup SnapMirror relationship.

You can use the `snapmirror create` command to create a data protection relationship or an extended data protection relationship between SnapLock source and SnapLock destination endpoints. When the cluster containing the source is running ONTAP 9.5.0 or later, the default relationship type is extended data protection (XDP), otherwise it is data protection (DP).

A SnapMirror relationship can also be created between a non-SnapLock source and a SnapLock destination to WORM-protect snapshots. In this case, the relationship type is always set to extended data protection (XDP).

The `snapmirror create` command can be used to create an extended data protection (XDP) relationship between a ONTAP volume and a non-ONTAP endpoint that supports SnapMirror (AltaVault, Solidfire). AltaVault endpoints can only be used as destinations. SolidFire endpoints can be used as sources or destinations.

The `snapmirror create` command can be used to create a synchronous relationship between FlexVol volumes, which provides zero Recovery Point Objective (RPO) data protection. SnapMirror Synchronous supports two policy types, *`sync-mirror`* and *`strict-sync-mirror`* . Upon a permanent replication failure, the strict-sync-mirror variant restricts further client IO on the primary, whereas the sync-mirror variant does not.

The `snapmirror create` command can be used to create a SnapMirror active sync relationship between Consistency Groups that have FlexVol volumes as items. This provides zero RPO and zero Recovery Time Objective (RTO) data protection, also known as transparent application failover, for SAN workloads. An SnapMirror active sync relationship supports the *`automated-failover`* or *`automated-failover-duplex`* policy type and can only be created when the systems containing the source items and the destination items are running ONTAP 9.9.1 or later.

Asynchronous SnapMirror relationships can also be created between the source Application Consistency Groups and the destination Application Consistency Groups. In this case, the relationship type is always set to extended data protection (XDP). An Application Consistency Group consists of FlexVol volumes as items and the SnapMirror relationships can only be created when the systems containing the source items and the destination items are running ONTAP 9.13.1 or later.

Before using this command to create a SnapMirror relationship between Vservers, you typically create a source and destination Vserver using the vserver create command. The source Vserver should be of *`subtypedefault`* and the destination Vserver of *`subtypedp-destination`* . Also, before creating the relationship between Vservers, you must setup Vserver peer by using the vserver peer create command between the source and destination Vservers. A Vserver relationship cannot be created between two Vservers within the same cluster. The root volume of the destination Vserver will be read-write and data from the source Vserver's root volume will not be replicated. Therefore there will be no volume level relationship created between the root volumes of the two Vservers.

After creating the relationship, the destination Vserver must be initialized by using the snapmirror initialize command.

Before using this command to create a volume SnapMirror relationship, you typically create a source and destination volume using the volume create command. The source volume should be in the online state and a read-write (RW) type. The destination volume should be in the online state and a data protection (DP) type. For FlexGroup SnapMirror relationships, the source and destination FlexGroup volumes must be spread over the same number of aggregates as specified in the `-aggr-list` parameter with the same number of constituents per aggregate as specified in the `-aggr-list-multiplier` parameter of the volume create command.

When a FlexGroup SnapMirror relationship is created, normally hidden relationships are also created for the constituent volumes. These relationships can be seen by using the `-expand` parameter of the snapmirror show command. Source information for these relationships can be seen using the `-expand` parameter of the snapmirror list-destinations command. Other SnapMirror commands are disabled for FlexGroup constituent relationships and FlexGroup constituent volumes.

If all systems involved are running ONTAP version 8.2 and later, a Vserver peering relationship must be set up using the vserver peer create command between the source and the destination Vservers to create a relationship between the source and destination volumes. To enable interoperability with ONTAP 8.1, if the source volume is on a storage system running ONTAP 8.1, the cluster administrator can create a data

protection relationship between the source and destination volumes without a Vserver peering relationship between the source and destination Vservers. These relationships are managed the same way as on ONTAP 8.1 and the *"Relationship Capability"* field, as shown in the output of the snapmirror show command, is set to *"Pre 8.2"*.

> SnapMirror relationships, except load-sharing relationships, which are created between two volumes which are both on a storage system running ONTAP version 8.2 and later have the *"Relationship Capability"* field set to *"8.2 and above"*.

Load-sharing mirrors must be confined to a single Vserver; they are not allowed to span Vservers. Load-sharing relationships are created with the *"Relationship Capability"* field set to *"Pre 8.2"* even if both the source and destination volumes are on a storage system running ONTAP version 8.2 and later. There is no *"8.2 and above"* implementation for load-sharing relationships.

A set of load-sharing mirrors can have one or more destination volumes. You create separate SnapMirror relationships between the common source volume and each destination volume to create the set of load-sharing mirrors.

The source or destination of a load-sharing relationship cannot be the source or destination of any other SnapMirror relationship.

After creating the relationship, the destination volume can be initialized using the snapmirror initialize command. The destination volumes in a set of load-sharing mirrors are initialized using the snapmirror initialize-ls-set command.

The `snapmirror create` command must be used from the destination Vserver or cluster.

Parameters

{ -S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name*.

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

-source-vserver <vserver name> - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not

supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name*.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

-destination-vserver <vserver name> - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

-destination-volume <volume name> - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

[-type <snapmirrorType>] - Relationship Type

This parameter specifies the type of SnapMirror relationship that will be created. You can create a data protection (DP) relationship (deprecated), an extended data protection (XDP) relationship or a load-sharing (LS) relationship. The default value is XDP for all relationships except for Vserver DR relationships. FlexGroup volumes, Application Consistency Groups and non-ONTAP endpoints support only XDP relationships. For FlexVol volume relationships, when DP is specified as the type, an XDP relationship will be created unless a relationship of type DP in the opposite direction already exists. In that case the new relationship will be of type DP.

[-vserver <vserver name>] - Managing Vserver

If this optional parameter is specified, it designates the managing Vserver. The managing Vserver is authorized to use SnapMirror commands to manage the SnapMirror relationship. The `-vserver` parameter is currently a reserved parameter.

[-schedule <text>] - SnapMirror Schedule

This optional parameter designates the name of the schedule which is used to update the SnapMirror relationship. If you do not designate a schedule, updates are not scheduled, so you must update the SnapMirror relationship manually using the snapmirror update command or, in the case of a set of load-sharing mirrors, using the snapmirror update-ls-set command.

> You define and name a schedule using the job schedule cron create command.

[-policy <sm_policy>] - SnapMirror Policy

This optional parameter designates the name of the SnapMirror policy which is associated with the SnapMirror relationship. For FlexVol volume relationships, the default policy when the data protection (DP) type is specified is *MirrorAllSnapshots* and the default policy when no type is specified or when the extended data protection (XDP) type is specified is *MirrorAndVault* . For FlexGroup volume and Application Consistency Group relationships, the *MirrorAndVault* policy is applied. For SnapMirror relationships between SnapLock volumes, if no policy is specified the default policy *MirrorAllSnapshots* is applied. For relationships with a SolidFire endpoint, there is no default policy. For these relationships a policy as described below must be specified. This parameter is not applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"* .

In ONTAP 8.2 data protection (DP) relationships were used for mirroring, while extended data protection (XDP) relationships were used for vaulting. In ONTAP 8.3 extended data protection (XDP) relationships support two more use cases, mirroring and unified mirror-vault. The exact behavior of an extended data protection (XDP) relationship is governed by the `snapmirror policy` associated with that relationship. In ONTAP 8.3 the `snapmirror policy` has a new field *type* to indicate how the relationships with that policy will behave. The supported types are *async-mirror* (mirroring), *vault* (vaulting) and *mirror-vault* (unified mirroring and vault). For XDP relationships between a ONTAP source volume and an AltaVault destination endpoint, only policies of type *vault* are supported. For XDP relationships between a ONTAP source volume and a SolidFire destination endpoint, only policies of type *async-mirror* without an *all_source_snapshots* rule are supported. For XDP relationships between a SolidFire source endpoint and a ONTAP destination volume, only policies of type *async-mirror* without an *all_source_snapshots* rule, and policies of type *mirror-vault* are supported. SnapMirror policies of type *async-mirror* associated with FlexVol volume relationships when relationship type DP is specified or when no relationship type is specified, must include the label *all_source_snapshots* . Refer to the man page for the snapmirror policy create command for more information.

> You define and name a policy using the snapmirror policy create command.

[-tries <integer_or_unlimited>] - Tries Limit

This optional parameter specifies the maximum number of times to attempt each manual or scheduled transfer for a SnapMirror relationship. The default is eight times. The `-tries` parameter can be set to *0* to disable manual and scheduled updates for the SnapMirror relationship. This parameter is only applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"* . For relationships with *"8.2 and above"* capability, the tries limit is controlled by the value of tries in the SnapMirror policy that is associated with the relationship.

[-k, -throttle <throttleType>] - Throttle (KB/sec)

This optional parameter limits the network bandwidth used for transfers. It configures for the relationship the maximum rate (in Kbytes/sec) at which data can be transferred. If no throttle is configured, by default the SnapMirror relationship fully utilizes the network bandwidth available. You can also configure the relationship to fully use the network bandwidth available by explicitly setting the throttle to *unlimited* or *0* . The minimum effective throttle value is four Kbytes/sec, so if you specify a throttle value between *1* and *4* ,

it will be treated as *4* . For FlexGroup volume relationships, the throttle value is applied individually to each constituent relationship. The `-throttle` parameter does not affect load-sharing mirrors and other SnapMirror relationships with *"Relationship Capability"* of *"Pre 8.2"* confined to a single cluster.

[-cg-item-mappings {<source volume>:@<destination volume>|<source item path>:@<destination item path>}] - Consistency Group Item Mappings

This optional parameter specifies a list of the consistency group (CG) item mappings. It is considered only if the supplied source and destination path values denote valid CG paths. For CG level relationships, this parameter must be specified. The value for this parameter must contain mappings of one or more pairs of constituent source and destination volumes of the form *srcvol1:@dstvol1,…* .

[-identity-preserve {true|false}] - Identity Preserve Vserver DR

Specifies whether or not the identity of the source Vserver is replicated to the destination Vserver of the Vserver SnapMirror relationship that will be created. If this parameter is set to true, the source Vserver's configuration will additionally be replicated to the destination. If the parameter is set to false, then only the source Vserver's volumes and RBAC configuration are replicated to the destination. This parameter is applicable only for SnapMirror relationships with Vserver endpoints. The default value is false.

[-is-auto-expand-enabled {true|false}] - Is Auto Expand Enabled

This optional parameter specifies whether or not a FlexGroup SnapMirror relationship and its destination FlexGroup volume should be auto-expanded if the source FlexGroup volume is expanded. This parameter is supported only for FlexGroup SnapMirror relationships. The default value is true. For Application Consistency Groups, this parameter is always set to true.

[-backoff-level {high|medium|none}] - SM Backoff Level due to Client Ops

This optional parameter specifies the SnapMirror backoff level due to client ops. This parameter is supported only for FlexVol SnapMirror relationships. The default value is `high` .

Examples

To create an extended data protection relationship between the source endpoint *vs1.example.com:dept_eng* , and the destination endpoint *vs2.example.com:dept_eng_dp_mirror2* , with the default policy of *MirrorAndVault* , type the following command:

```
vs2.example.com::> snapmirror create -destination-path
            vs2.example.com:dept_eng_dp_mirror2 -source-path
            vs1.example.com:dept_eng
```

To create an extended data protection relationship between the source FlexGroup *vs1.example.com:fg_src* and the destination FlexGroup *vs2.example.com:fg_dst* , with the default policy of *MirrorAndVault* , type the following command:

```
vs2.example.com::> snapmirror create -destination-path
            vs2.example.com:fg_dst -source-path
            vs1.example.com:fg_src
```

To create a SnapMirror Synchronous relationship between the source FlexVol *vs1.example.com:vol_log* , and the destination FlexVol *vs2.example.com:vol_log_sync_dp* when the source cluster is running

ONTAP 9.5 or above, type the following command:

```
vs2.example.com::> snapmirror create -destination-path
            vs2.example.com:vol_log_sync_dp -source-path
            vs1.example.com:vol_log -policy Sync
```

To create a strict SnapMirror Synchronous relationship between the source FlexVol *vs1.example.com:vol_log* , and the destination FlexVol *vs2.example.com:vol_log_sync_dp* when the source cluster is running ONTAP 9.5 or above, type the following command:

```
vs2.example.com::> snapmirror create -destination-path
            vs2.example.com:vol_log_sync_dp -source-path
            vs1.example.com:vol_log -policy StrictSync
```

To create a data protection mirror between the source endpoint *cluster1://vs1.example.com/dept_eng* , and the destination endpoint *cluster2://vs2.example.com/dept_eng_dp_mirror2* when the source cluster is running ONTAP 8.1 software, type the following command:

```
cluster2::> snapmirror create -destination-path
           cluster2://vs2.example.com/dept_eng_dp_mirror2 -source-path
           cluster1://vs1.example.com/dept_eng
    -type DP
```

To create a load-sharing mirror between the source endpoint *cluster1://vs1.example.com/vs1_root* which is a Vserver root volume, and the destination endpoint *cluster1://vs1.example.com/vs1_root_ls1* with the schedule named *5min* used to update the relationship, type the following command:

```
cluster1::> snapmirror create
       -destination-path cluster1://vs1.example.com/vs1_root_ls1
       -source-path      cluster1://vs1.example.com/vs1_root -type LS
-schedule 5min
```

To create a SnapMirror relationship between the source Vserver *vs1.example.com* , and the destination Vserver *dvs1.example.com* with the schedule named *hourly* used to update the relationship, type the following command:

```
cluster2::> snapmirror create
       -destination-path dvs1.example.com:
       -source-path vs1.example.com:
       -schedule hourly
```

To create an extended data protection (XDP) relationship between the ONTAP source endpoint *vs1.example.com:data_ontap_vol* , and the AltaVault destination endpoint *10.0.0.11:/share/share1* , type the following command:

```
vs1.example.com::> snapmirror create -destination-path
      10.0.0.11:/share/share1 -source-path
      vs1.example.com:data_ontap_vol
      -type XDP
```

To create an extended data protection (XDP) relationship between the SolidFire source endpoint *10.0.0.12:/lun/0001* , and the ONTAP destination endpoint *vs2.example.com:data_ontap_vol2* , type the following command:

```
vs2.example.com::> snapmirror create -source-path
      10.0.0.12:/lun/0001 -destination-path
      vs2.example.com:data_ontap_vol2
      -type XDP -policy MirrorLatest
```

To create an SnapMirror active sync relationship with the following attributes:

- It is between the source Consistency Group *cg_src* in *Vserver vs1.example.com* , and the destination Consistency Group *cg_dst* in Vserver *vs2.example.com* .
- It has item mappings between volumes *srcvol1* and *srcvol2* and volumes *dstvol1* and *dstvol2* .
- It can use the default policy named *AutomatedFailOver* that has a policy type of *automated-failover* or *AutomatedFailOverDuplex* that has a policy type of *automated-failover-duplex* .

Type the following command from the destination cluster:

```
destination::> snapmirror create -destination-path
        vs2.example.com:/cg/cg_dst -source-path
        vs1.example.com:/cg/cg_src
  -policy AutomatedFailOver
        -cg-item-mappings srcvol1:@dstvol1,srcvol2:@dstvol2
```

To create an extended data protection relationship between the source Application Consistency Group *vs1.example.com:/cg/cg_src* and the destination Application Consistency Group *vs2.example.com:/cg/cg_dst* , with the policy of *MirrorAllSnapshots* , type the following command:

```
destination::> snapmirror create -destination-path
        vs2.example.com:/cg/cg_dst -source-path
  vs1.example.com:/cg/cg_src
  -policy MirrorAllSnapshots
        -cg-item-mappings srcvol1:@dstvol1,srcvol2:@dstvol2
```

Related Links

- vserver create
- vserver peer create
- snapmirror initialize
- volume create
- snapmirror show
- snapmirror list-destinations
- snapmirror initialize-ls-set
- snapmirror update
- snapmirror update-ls-set
- job schedule cron create
- snapmirror policy create

snapmirror delete

Delete a SnapMirror relationship

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror delete` command removes the SnapMirror relationship between a source endpoint and a destination endpoint. The destination endpoint can be a Vserver, a volume, an Application Consistency Group, or a non-ONTAP endpoint. The Vservers, volumes, FlexVol volumes of the Consistency Group and non-ONTAP destinations are not destroyed and snapshots on the volumes are not removed.

For relationships with SolidFire endpoints, the SnapMirror source commands snapmirror release and snapmirror list-destinations are not supported. Therefore, snapshots that are locked by SnapMirror on the source container cannot be cleaned up using the snapmirror release command. If the source container is a ONTAP volume, in order to reclaim space captured in the base snapshot on the volume, issue a `snapshot delete` command specifying the `-ignore-owners` parameter in diag privilege level. To reclaim space captured in a snapshot locked by SnapMirror on a SolidFire system, refer to SolidFire documentation.

The `snapmirror delete` command fails if a SnapMirror transfer for the SnapMirror relationship is in progress for relationships with *"Relationship Capability"* of *"Pre 8.2"*. For relationships with *"8.2 and above"* capability the delete will succeed even if a transfer is in progress and the transfer will ultimately stop.

A set of load-sharing mirrors can contain multiple destination volumes, each destination volume having a separate SnapMirror relationship with the common source volume. When used on one of the SnapMirror relationships from the set of load-sharing mirrors, the `snapmirror delete` command deletes the specified SnapMirror relationship from the set of load-sharing mirrors.

The `snapmirror delete` command preserves the read-write or read-only attributes of the volumes of a SnapMirror relationship after the relationship is deleted. Therefore, a read-write volume that was the source of a SnapMirror relationship retains its read-write attributes, and a data protection volume or a load-sharing volume that was a destination of a SnapMirror relationship retains its read-only attributes. Similarly, the

subtype attribute of source and destination Vservers is not modified when a Vserver SnapMirror relationship is deleted.

> When a SnapMirror relationship from a set of load-sharing mirrors is deleted, the destination volume becomes a data protection volume and retains the read-only attributes of a data protection volume.

For relationships with a policy of type *strict-sync-mirror* or *sync-mirror* , the relationship must be *Quiesced* before it can be deleted.

This command is supported for SnapMirror relationships with the field *"Relationship Capability"* showing as either *"8.2 and above"* or *"Pre 8.2"* in the output of the snapmirror show command.

For relationships with *"Relationship Capability"* of *"8.2 and above"* , the `snapmirror delete` command must be used from the destination Vserver or cluster. The SnapMirror relationship information is deleted from the destination Vserver, but no cleanup or deletion is performed on the source Vserver. The snapmirror release command must be issued on the source Vserver to delete the source relationship information.

For relationships with *"Relationship Capability"* of *"Pre 8.2"* , you can use this command from the source or from the destination cluster. When used from the destination cluster, the SnapMirror relationship information on the source and destination clusters is deleted. When used from the source cluster, only the SnapMirror relationship information on the source cluster is deleted.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"* , a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* . For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name* . For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name* .

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"* . This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* .

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"* , `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name*.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

-destination-vserver <vserver name> - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

-destination-volume <volume name> - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

[-f, -force <true>] - Force

If specified, the delete proceeds even if it cannot clean up all artifacts of the relationship.

[-w, -foreground <true>] - Foreground Process

This specifies whether the operation runs as a foreground process. If this parameter is specified, the default setting is *true* (the operation runs in the foreground). When set to *true*, the command will not return until the process completes. This parameter is only applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"*.

Examples

To delete the SnapMirror relationship with the destination endpoint *vs2.example.com:dept_eng_dp_mirror4*, type the following command:

```
vs2.example.com::> snapmirror delete -destination-path
       vs2.example.com:dept_eng_dp_mirror4
```

For relationships with *"Relationship Capability"* of *"Pre 8.2"* , to delete the SnapMirror relationship with the destination endpoint *cluster2://vs2.example.com/dept_eng_dp_mirror4* , type the following command:

```
cluster2::> snapmirror delete -destination-path
         cluster2://vs2.example.com/dept_eng_dp_mirror4
```

To delete the SnapMirror relationship with destination endpoint *dvs1.example.com:* , type the following command:

```
cluster2::> snapmirror delete -destination-path
   dvs1.example.com:
```

To delete the SnapMirror active sync relationship with the destination Consistency Group *cg_dst* in Vserver *vs2.example.com* , type the following command from the destination cluster:

```
destination::> snapmirror delete -destination-path
        vs2.example.com:/cg/cg_dst
```

To delete the Application Consistency Group SnapMirror relationship with the destination Application Consistency Group *app_cg_dst* in Vserver *vs2.example.com* and all item mappings, type the following command:

```
vs2.example.com::> snapmirror delete -destination-path
            vs2.example.com:/cg/app_cg_dst
```

Related Links

- snapmirror release
- snapmirror list-destinations
- snapmirror show

snapmirror flexgroup-epuuid-prefix-prepare-to-downgrade

Prepares the system for downgrade

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `flexgroup-epuuid-prefix-prepare-to-downgrade` command prepares nodes to downgrade to a release without FlexGroup endpoint prefixed with the group endpoint uuid capability. Prior to disabling the capability, remove all FlexGroup SnapMirror object store relationships which have a source or destination endpoint prefixed with the group endpoint uuid.

Examples

The following example disables the FlexGroup endpoint prefixed with group endpoint uuid capability in the local cluster:

```
cluster1::> snapmirror flexgroup-epuuid-prefix-prepare-to-downgrade
```

snapmirror initialize-ls-set

Start a baseline load-sharing set transfer

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror initialize-ls-set` command initializes and updates a set of load-sharing mirrors. This command is usually used after the snapmirror create command is used to create a SnapMirror relationship for each of the destination volumes in the set of load-sharing mirrors. The initial transfers to empty load-sharing mirrors are baseline transfers done in parallel. During a baseline transfer ONTAP takes a snapshot on the source volume to capture the current image of the source volume and transfers all of the snapshots on the source volume to each of the destination volumes.

After the `snapmirror initialize-ls-set` command successfully completes, the last snapshot transferred is made the exported snapshot on the destination volumes.

The parameter that identifies the set of load-sharing mirrors is the source volume. Data and snapshots are transferred from the source volume to all up-to-date destination volumes in the set of load-sharing mirrors.

Use the snapmirror initialize command to add and initialize a new destination volume to an existing set of load-sharing mirrors.

> Even if the load-sharing set only has one mirror, you still need to use the `snapmirror initialize-ls-set` command to initialize the set. The snapmirror initialize command can only be used to initialize a new destination volume, if the load-sharing set has already been initialized.

This command is only supported for SnapMirror relationships with the field *`"Relationship Capability"`* showing as *`"Pre 8.2"`* in the output of the snapmirror show command.

Parameters

{ -S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name*.

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

-source-vserver <vserver name> - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

-source-volume <volume name> - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-w, -foreground <true>] - Foreground Process

This specifies whether the operation runs as a foreground process. If this parameter is specified, the default setting is *true* (the operation runs in the foreground). When set to *true*, the command will not return until the process completes. This parameter is only applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"*.

Examples

To initialize the group of load-sharing mirrors for the source endpoint *//vs1.example.com/vs1_root*, type the following command:

```
cluster1::> snapmirror initialize-ls-set -source-path
       //vs1.example.com/vs1_root
```

Related Links

- snapmirror create

- snapmirror initialize
- snapmirror show

snapmirror initialize

Start a baseline transfer

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror initialize` command initializes the destination Vserver, volume, Application Consistency Group or a non-ONTAP endpoint of a SnapMirror relationship. The command behaves differently between data protection (DP), extended data protection (XDP) and load-sharing (LS) relationships.

If you specify a *`sync-mirror`* or *`strict-sync-mirror`* type policy, the `snapmirror initialize` command creates and initializes a synchronous relationship and brings it InSync, providing zero RPO data protection.

If you are initializing a relationship with a policy of type *`automated-failover`* the `snapmirror initialize` command initializes a SnapMirror active sync relationship and brings the status to InSync, providing zero RPO and zero RTO data protection.

If you are initializing a relationship with a policy of type *`automated-failover-duplex`* the `snapmirror initialize` command initializes a SnapMirror active sync relationship and brings the status to InSync, providing zero RPO, zero RTO data protection and load balanced solution.

SnapMirror active sync is only supported for SAN. As a result, NAS access is not supported on SnapMirror active sync volumes and SnapMirror active sync initialization is not supported on volumes with NAS access.

For data protection (DP) and extended data protection (XDP) relationships, the `snapmirror initialize` command initializes the destination volume.

For load-sharing (LS) relationships, the `snapmirror initialize` command initializes a new load-sharing mirror in an existing set of load-sharing mirrors. If the command finishes before the start of a scheduled or manual transfer of the set of load-sharing mirrors, the load-sharing mirror is up to date with the set of load-sharing mirrors; otherwise, the load-sharing mirror will be brought up to date at the next scheduled or manual transfer of the set of load-sharing mirrors.

The initial transfer to an empty destination volume is called a baseline transfer. During a baseline transfer for a data protection (DP) or extended data protection (XDP) relationship, the `snapmirror initialize` command takes a snapshot on the source volume to capture the current image of the source volume. For data protection relationships, the `snapmirror initialize` command transfers all of the snapshots up to and including the snapshot created by it from the source volume to the destination volume. For extended data protection (XDP) relationships, the `snapmirror initialize` command behavior depends on the `snapmirror policy` associated with the relationship. If the policy type is *`async-mirror`* then depending on the rules in the policy it can transfer either all the snapshots up to and including the snapshot created by it or only the snapshot created by it from the source volume to the destination volume. For extended data protection (XDP) relationships with policy type *`vault`* or *`mirror-vault`* the `snapmirror initialize` transfers only the snapshot created by it.

After the `snapmirror initialize` command successfully completes, the last snapshot transferred is made the exported snapshot on the destination volume.

You can use the `snapmirror initialize` command to initialize a specific load-sharing mirror that is new to the set of load-sharing mirrors. An initialize of the new load-sharing mirror should bring it up to date with the other up-to-date destination volumes in the set of load-sharing mirrors.

> Using the `snapmirror initialize` command to initialize a set of load-sharing mirrors will not work. Use the snapmirror initialize-ls-set command to initialize a set of load-sharing mirrors.

If a SnapMirror relationship does not already exist, that is, the relationship was not created using the snapmirror create command, the `snapmirror initialize` command will implicitly create the SnapMirror relationship, with the same behaviors as described for the snapmirror create command before initializing the relationship. This implicit create feature is not supported for Vservers.

This command is supported for SnapMirror relationships with the field *"Relationship Capability"* showing as either *"8.2 and above"* or *"Pre 8.2"* in the output of the snapmirror show command.

For relationships with *"Relationship Capability"* of *"8.2 and above"* , you can track the progress of the operation using the snapmirror show command.

For relationships with *"Relationship Capability"* of *"Pre 8.2"* , a job will be spawned to operate on the SnapMirror relationship, and the job id will be shown in the command output. The progress of the job can be tracked using the job show and job history show commands.

The `snapmirror initialize` command must be used from the destination Vserver or cluster.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"* , a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* . For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name* . For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name* .

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"* . This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* .

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"* , `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name*.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

-destination-vserver <vserver name> - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

-destination-volume <volume name> - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

[-s, -source-snapshot <text>] - Source Snapshot

This optional parameter specifies the snapshot that `snapmirror initialize` will use for the baseline transfer. For data protection (DP) relationships, the baseline transfer will include all of the snapshots up to and including the specified snapshot. For extended data protection (XDP) relationships, the `snapmirror initialize` command behavior depends on the `snapmirror policy` associated with the relationship. If the policy type is *async-mirror* then depending on the rules in the policy it can transfer either all the snapshots up to and including the specified snapshot or only the specified snapshot from the source volume to the destination volume. For extended data protection (XDP) relationships with policy type *vault* or *mirror-vault* the `snapmirror initialize` transfers only the specified snapshot. This parameter is not supported for SnapMirror Synchronous relationships and relationships with *"Relationship Capability"* of *"Pre 8.2"*.

[-type <snapmirrorType>] - Snapmirror Relationship Type

Specifies the type of SnapMirror relationship if a relationship is implicitly created. This parameter is the same as the one used in the snapmirror create command.

[-policy <sm_policy>] - SnapMirror Policy

This optional parameter designates the name of the SnapMirror policy which is associated with the SnapMirror relationship. If you do not designate a policy, the current policy will be retained. This parameter is not applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"*.

> You define and name a policy using the snapmirror policy create command.

[-k, -throttle <throttleType>] - Throttle (KB/sec)

This optional parameter limits the network bandwidth used for the initialize transfer. It sets the maximum rate (in Kbytes/sec) at which data can be transferred during the operation. If this parameter is not specified, the throttle value configured for the relationship with the snapmirror create or snapmirror modify command will be used. To fully use the network bandwidth available, set the throttle value to *unlimited* or *0*. The minimum throttle value is four Kbytes/sec, so if you specify a throttle value between *1* and *4*, it will be treated as if you specified *4*. For FlexGroup relationships, the throttle value is applied individually to each constituent relationship. For SnapMirror Synchronous relationships the throttle value is applicable only for asynchronous transfers performed as part of user-initiated operations. The `-throttle` parameter does not affect load-sharing transfers and transfers for other relationships with *"Relationship Capability"* of *"Pre 8.2"* confined to a single cluster.

[-transfer-priority {low|normal}] - Transfer Priority

This optional parameter specifies the priority at which the transfer runs. The default value for this parameter is the value in the SnapMirror policy associated with this relationship. This parameter is not applicable to relationships with a *"Relationship Capability"* of *"Pre 8.2"*.

[-cg-item-mappings {<source volume>:@<destination volume>|<source item path>:@<destination item path>}] - Consistency Group Item Mappings

This optional parameter specifies a list of the consistency group (CG) item mappings. It is considered only if the supplied source and destination path values denote valid CG paths. For CG level relationships, this parameter must be specified. The value for this parameter must contain mappings of one or more pairs of constituent source and destination volumes of the form *srcvol1:@dstvol1,…*.

[-is-auto-expand-enabled {true|false}] - Is Auto Expand Enabled

This optional parameter specifies whether or not a FlexGroup SnapMirror relationship and its destination FlexGroup should be auto-expanded if the source FlexGroup is expanded. This parameter is supported only for FlexGroup SnapMirror relationships. If this initialize is creating a new SnapMirror relationship, the default value is true. If it is not creating a new relationship, if a value is specified, it must match the current value for the existing relationship. If the parameter is not specified the existing value will be retained.

[-backoff-level {high|medium|none}] - SM Backoff Level due to Client Ops

This optional parameter specifies the SnapMirror backoff level due to client ops. This parameter is supported only for FlexVol SnapMirror relationships. The default value is `high`.

[-w, -foreground <true>] - Foreground Process

This specifies whether the operation runs as a foreground process. If this parameter is specified, the default setting is *true* (the operation runs in the foreground). When set to *true*, the command will not return until the process completes. This parameter is only applicable to relationships with *"Relationship*

Capability" of *"Pre 8.2"*.

Examples

To start the initial transfer for the SnapMirror relationship with the destination endpoint *vs2.example.com:dept_eng_dp_mirror2* after the relationship has been created with the snapmirror create command, type the following command:

```
vs2.example.com::> snapmirror initialize -destination-path
       vs2.example.com:dept_eng_dp_mirror2
```

For relationships with *"Relationship Capability"* of *"Pre 8.2"*, to start the initial transfer for the SnapMirror relationship with the destination endpoint *cluster2://vs2.example.com/dept_eng_dp_mirror2* after the relationship has been created with the snapmirror create command, type the following command:

```
cluster2::> snapmirror initialize -destination-path
         cluster2://vs2.example.com/dept_eng_dp_mirror2
```

To create a data protection mirror relationship between the source endpoint *vs1.example.com:dept_mkt*, and the destination endpoint *vs2.example.com:dep_mkt_dp_mirror*, and start the initial transfer, type the following command:

```
vs2.example.com::> snapmirror initialize -destination-path
         vs.example.com2:dept_mkt_dp_mirror
         -source-path vs1.example.com:dept_mkt
```

To create a data protection mirror relationship between the source endpoint *cluster1://vs1.example.com/dept_mkt*, and the destination endpoint *cluster2://vs2.example.com/dep_mkt_dp_mirror*, and start the initial transfer when the source cluster is running ONTAP 8.1 software, type the following command:

```
cluster2::> snapmirror initialize -destination-path
         cluster2://vs2.example.com/dept_mkt_dp_mirror
         -source-path cluster1://vs1.example.com/dept_mkt
```

To create an extended data protection (XDP) relationship between the ONTAP source endpoint *vs1.example.com:data_ontap_vol*, and the AltaVault destination endpoint *10.0.0.11:/share/share1*, and start the initial transfer, type the following command:

```
vs1.example.com::> snapmirror initialize -destination-path
  10.0.0.11:/share/share1
  -source-path vs1.example.com:data_ontap_vol -type XDP
```

To start the initial transfer for the Vserver SnapMirror relationship with destination endpoint *dvs1.example.com:* after the relationship was created with the snapmirror create command, type the following command:

```
cluster2::> snapmirror initialize -destination-path
         dvs1.example.com:
```

To initialize the SnapMirror Synchronous relationship between FlexVol volumes *vol_log* and *vol_log_sync_dp* and bring it to InSync, after it is created using the snapmirror create command, type the following command:

```
vs2.example.com::> snapmirror initialize -destination-path
         vs2.example.com:vol_log_sync_dp
```

To create a strict SnapMirror Synchronous relationship between FlexVol volumes *vol_log* and *vol_log_sync_dp* , to initialize it and bring it to InSync, type the following command:

```
vs2.example.com::> snapmirror initialize -destination-path
         vs2.example.com:vol_log_sync_dp -source-path
         vs1.example.com:vol_log -policy StrictSync
```

To create and initialize an SnapMirror active sync relationship with the following attributes:

- It is between the source Consistency Group *cg_src* in Vserver *vs1.example.com* , and the destination Consistency Group *cg_dst* in Vserver *vs2.example.com* .
- It has item mappings between volumes *srcvol1* and *srcvol2* and volumes *dstvol1* and *dstvol2* .
- It uses the policy named *AutomatedFailOver* that has a policy type of *automated-failover* .

Type the following command from the destination cluster:

```
destination::> snapmirror initialize -destination-path
        vs2.example.com:/cg/cg_dst -source-path
  vs1.example.com:/cg/cg_src
  -policy AutomatedFailOver
        -cg-item-mappings srcvol1:@dstvol1,srcvol2:@dstvol2
```

To initialize the previously created SnapMirror active sync relationship, type the following command from the destination cluster:

```
destination::> snapmirror initialize -destination-path
        vs2.example.com:/cg/cg_dst
```

To create and initialize an Application Consistency Group relationship with the following attributes:

- It is between the source Application Consistency Group *app_cg_src* in Vserver *vs1.example.com* , and the destination Application Consistency Group *app_cg_dst* in Vserver *vs2.example.com* . It has item mappings between volumes *srcvol1* and *dstvol2* and volumes *srcvol2* and *dstvol2* . It uses the default policy named *MirrorAndVault* that has a policy type of *mirror-vault* .

Type the following command from the destination cluster:

```
destination::> snapmirror initialize -destination-path
        vs2.example.com:/cg/app_cg_dst -source-path
        vs1.example.com:/cg/app_cg_src
        -cg-item-mappings srcvol1:@dstvol1,srcvol2:@dstvol2
```

To initialize the previously created Application Consistency Group relationship, type the following command from the destination cluster:

```
destination::> snapmirror initialize -destination-path
        vs2.example.com:/cg/app_cg_dst
```

Related Links

- snapmirror initialize-ls-set
- snapmirror create
- snapmirror show
- job show
- job history show
- snapmirror policy create
- snapmirror modify

snapmirror list-destinations

Display a list of destinations for SnapMirror sources

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror list-destinations` command displays information including the destination endpoints, the relationship status, and transfer progress, for SnapMirror relationships whose source endpoints are in the current Vserver if you are in a Vserver context, or the current cluster if you are in a cluster context. +

The command might display several relationships that have the same source and destination endpoints, but have different relationship IDs. If this is the case, some of the information is stale. It corresponds to relationships that have been deleted on the destination Vserver or cluster, and have not been released yet on the source Vserver or source cluster.

The relationships and the information displayed are controlled by the parameters that you specify. If no parameters are specified, the command displays the following information associated with each SnapMirror relationship whose source endpoint is in the current Vserver if you are in a Vserver context, or the current cluster if you are in a cluster context:

- Source path
- Relationship Type
- Destination Path
- Relationship Status
- Transfer Progress
- Progress Last Updated
- Relationship ID

Note the following limitations on the information displayed by the `snapmirror list-destinations` command:

- The *"Relationship Status"* field is not valid after the node hosting the source volume joins the cluster quorum, until at least one transfer is performed on the SnapMirror relationship.
- *"Transfer Progress"* and *"Progress Last Updated"* fields are only valid if a snapshot transfer is in progress.
- The *"Relationship ID"* field is not valid for Vserver SnapMirror relationships.
- The *"Relationship Status"*, *"Transfer Progress"*, and *"Progress Last Updated"* fields are not valid for FlexGroup relationships, but they are valid for FlexGroup constituent relationships.

The `-instance` and `-fields` parameters are mutually exclusive and select the fields that are displayed. The `-instance` parameter if specified, displays detailed information about the relationships. The `-fields` parameter specifies what fields should be displayed. The other parameters of the `snapmirror list-destinations` command, select the SnapMirror relationships for which the information is displayed.

This command is not supported for SnapMirror relationships with non-ONTAP endpoints.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>,…` parameter, the command only displays the fields that you have specified.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all relationships selected.

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

Selects SnapMirror relationships that have a matching source path name.

| [-source-vserver <vserver name>] - Source Vserver

Selects SnapMirror relationships that have a matching source Vserver name.

[-source-volume <volume name>] - Source Volume }

Selects SnapMirror relationships that have a matching source volume name.

{ [-destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Destination Path

Selects SnapMirror relationships that have a matching destination path name.

| [-destination-vserver <vserver name>] - Destination Vserver

Selects SnapMirror relationships that have a matching destination Vserver name.

[-destination-volume <volume name>] - Destination Volume }

Selects SnapMirror relationships that have a matching destination volume name.

[-relationship-id <UUID>] - Relationship ID

Selects SnapMirror relationships that have a matching relationship identifier. This parameter is not supported for Vserver SnapMirror relationships.

[-type <snapmirrorType>] - Relationship Type

Selects SnapMirror relationships that have a matching relationship type. Possible values are:

- DP
- XDP
- RST

[-relationship-group-type {none|vserver|consistencygroup|flexgroup|vserverflexgroup|vserverconsistencygroup}] - Relationship Group Type

Selects SnapMirror relationships that have a matching relationship group type. Possible values are:

- *none*
- *vserver*
- *flexgroup*

[-policy-type {vault|async-mirror|mirror-vault|strict-sync-mirror|sync-mirror|automated-failover|automated-failover-duplex|continuous}] - SnapMirror Policy Type

Selects SnapMirror relationships that have a matching SnapMirror policy type. Possible values are:

- async-mirror
- vault
- mirror-vault

[-cg-item-mappings {<source volume>:@<destination volume>|<source item path>:@<destination item path>}] - Consistency Group Item Mappings

Selects SnapMirror relationships that have a matching list of Consistency Group (CG) item mappings. This parameter is applicable to CG level relationships only.

[-status <mirror status>] - Relationship Status

Selects SnapMirror relationships that have a matching relationship status. Possible values are:

- Idle
- Transferring

This parameter is not supported for FlexGroup and Application Consistency Group SnapMirror relationships, but it is supported for FlexGroup and Application Consistency Group constituent relationships.

[-transfer-progress {<integer>[KB|MB|GB|TB|PB]}] - Transfer Progress

Selects SnapMirror relationships that have a matching transfer progress. This parameter is not supported for FlexGroup and Application Consistency Group SnapMirror relationships, but it is supported for FlexGroup and Application Consistency Group constituent relationships.

[-progress-last-updated <MM/DD HH:MM:SS>] - Timestamp of Last Progress Update

Selects SnapMirror relationships that have a matching transfer progress last updated timestamp. This parameter is not supported for FlexGroup and Application Consistency Group SnapMirror relationships, but it is supported for FlexGroup and Application Consistency Group constituent relationships.

[-source-volume-node <nodename>] - Source Volume Node Name

Selects SnapMirror relationships that have a matching source volume node name. For FlexGroup relationships, it is the node which owns the root constituent source volume. This parameter is not supported for Vserver SnapMirror relationships.

[-expand <true>] - Show Constituents of the Group

Specifies whether to display constituent relationships of Vserver and FlexGroup SnapMirror relationships. By default, the constituents are not displayed.

Examples

To display summary information about all relationships whose source endpoints are in the current cluster, type the following command:

```
cluster1::> snapmirror list-destinations
Progress
Source             Destination          Transfer  Last      Relationship
Path         Type  Path         Status  Progress  Updated   ID
----------- ----- ------------ ------- --------- ---------
----------------
vserver1.example.com:dp_s1
            DP    vserver2.example.com:dp_d1
                                Idle    -         -         06b4327b-954f-
11e1-af65-123478563412
vserver1.example.com:xdp_s1
            XDP   vserver2.example.com:xdp_d1
                                Idle    -         -         a9c1db0b-954f-
11e1-af65-123478563412
vserver2.example.com:
            DP    dvserver2.example2.com:
                                Idle    -         -         -
3 entries were displayed.
```

To display summary information about all relationships whose source endpoints are in the current Vserver, type the following command:

```
vserver1.example.com::> snapmirror list-destinations
Progress
Source             Destination          Transfer  Last      Relationship
Path         Type  Path         Status  Progress  Updated   ID
----------- ----- ------------ ------- --------- ---------
----------------
vserver1.example.com:dp_s1
            DP    vserver2.example.com:dp_d1
                                Idle    -         -         06b4327b-954f-
11e1-af65-123478563412
vserver1.example.com:xdp_s1
            XDP   vserver2.example.com:xdp_d1
                                Idle    -         -         a9c1db0b-954f-
11e1-af65-123478563412
2 entries were displayed.
```

To display detailed information about SnapMirror relationships whose source endpoints are in the current Vserver, type the following command:

```
vserver1.example.com::> snapmirror list-destinations -instance
Source Path: vserver1.example.com:dp_s1
         Destination Path: vserver2.example.com:dp_d1
        Relationship Type: DP
  Relationship Group Type: none
      Relationship Status: Idle
        Transfer Progress: -
    Progress Last Updated: -
       Source Volume Node: node1
          Relationship ID: 06b4327b-954f-11e1-af65-123478563412
Source Path: vserver1.example.com:xdp_s1
         Destination Path: vserver2.example.com:xdp_d1
        Relationship Type: XDP
  Relationship Group Type: none
      Relationship Status: Idle
        Transfer Progress: -
    Progress Last Updated: -
       Source Volume Node: node2
          Relationship ID: a9c1db0b-954f-11e1-af65-123478563412

2 entries were displayed.
```

To display summary information about all relationships including constituent relationships whose source endpoints are in the current Vserver, type the following command:

```
cluster-1::> snapmirror list-destinations -expand
                                                   Progress
Source             Destination         Transfer   Last
Relationship
Path         Type  Path         Status Progress   Updated      Id
----------- ----- ------------ ------- --------- ------------
---------------
vs1:fg_s1   XDP   vs1:fg_d1    -       -         -            504abc00-
70a8-11e6-82be-0050568536d7
vs1:fg_s1__0001
            XDP   vs1:fg_d1__0001
                                -       -         -            5041f2aa-
70a8-11e6-82be-0050568536d7
vs1:fg_s1__0002
            XDP   vs1:fg_d1__0002
                                -       -         -            50421733-
70a8-11e6-82be-0050568536d7
vs1:fg_s1__0003
            XDP   vs1:fg_d1__0003
                                -       -         -            50421826-
70a8-11e6-82be-0050568536d7
vs1:fg_s1__0004
            XDP   vs1:fg_d1__0004
                                -       -         -            504218f0-
70a8-11e6-82be-0050568536d7
5 entries were displayed.
```

To display information about Application Consistency Group SnapMirror relationships type the following command:

```
cluster-1::> snapmirror list-destinations
                                                   Progress
Source             Destination         Transfer   Last
Relationship
Path         Type  Path         Status Progress   Updated      Id
----------- ----- ------------ ------- --------- ------------
---------------
vs0:/cg/cgA XDP   vs1:/cg/cg_dst_async Idle -    -            8cb8bc1a-
c193-11ed-9f4c-005056a7c7a2
```

snapmirror modify

Modify a SnapMirror relationship

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror modify` command allows you to change one or more properties of SnapMirror relationships. The key parameter that identifies any SnapMirror relationship is the destination endpoint. The destination endpoint can be a Vserver, a volume, an Application Consistency Group or a non-ONTAP endpoint.

For load-sharing mirrors, a change to a property affects all of the SnapMirror relationships in the set of load-sharing mirrors. Destination volumes in a set of load-sharing mirrors do not have individual property settings.

Changes made by the `snapmirror modify` command do not take effect until the next manual or scheduled update of the SnapMirror relationship. Changes do not affect updates that have started and have not finished yet.

This command is supported for SnapMirror relationships with the field *"Relationship Capability"* showing as either *"8.2 and above"* or *"Pre 8.2"* in the output of the snapmirror show command.

The `snapmirror modify` command must be used from the destination Vserver or cluster.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"* , a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* . For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name* . For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name* .

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"* . This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* .

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"* , `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"* , `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"* , a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* . For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name* .

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"* . This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"* .

-destination-vserver <vserver name> - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"* , `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

-destination-volume <volume name> - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"* , `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

[-vserver <vserver name>] - Managing Vserver

If this optional parameter is specified, designates the managing Vserver. The managing Vserver is authorized to use some SnapMirror commands to manage the SnapMirror relationship. The `-vserver` option is currently a reserved option.

[-schedule <text>] - SnapMirror Schedule

This optional parameter designates the name of the schedule which is used to update the SnapMirror relationship. If you do not designate a schedule, updates are not scheduled, so you must update the SnapMirror relationship manually using the snapmirror update command or, in the case of a set of load-sharing mirrors, using the snapmirror update-ls-set command.

> You define and name a schedule using the job schedule cron create command.

[-policy <sm_policy>] - SnapMirror Policy

This optional parameter designates the name of the SnapMirror policy which is associated with the SnapMirror relationship. If you do not designate a policy, the current policy will be retained. Modification of the policy is not supported for relationships with a policy of type *strict-sync-mirror* , *sync-mirror* , or *automated-failover-duplex* . If you want to use a different policy for SnapMirror Synchronous relationships, you need to delete the relationship and create it again with the new policy. This parameter is not applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"* .

You define and name a policy using the snapmirror policy create command.

[-tries <integer_or_unlimited>] - Tries Limit

This optional parameter specifies the maximum number of times to attempt each manual or scheduled transfer for a SnapMirror relationship. The default is eight times. The `-tries` parameter can be set to *0* to disable manual and scheduled updates for the SnapMirror relationship. This parameter is only applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"* . For relationships with *"8.2 and above"* capability, the tries limit is controlled by the value of tries in the SnapMirror policy that is associated with the relationship.

[-k, -throttle <throttleType>] - Throttle (KB/sec)

This optional parameter limits the network bandwidth used for transfers. It configures for the relationship the maximum rate (in Kbytes/sec) at which data can be transferred. If no throttle is configured, by default the SnapMirror relationship fully utilizes the network bandwidth available. You can also configure the relationship to fully use the network bandwidth available by explicitly setting the throttle to *unlimited* or *0* . The minimum effective throttle value is four Kbytes/sec, so if you specify a throttle value between *1* and *4* , it will be treated as *4* . For FlexGroup volume relationships, the throttle value is applied individually to each constituent relationship. The `-throttle` parameter does not affect load-sharing mirrors and other SnapMirror relationships with *"Relationship Capability"* of *"Pre 8.2"* confined to a single cluster.

[-is-auto-expand-enabled {true|false}] - Is Auto Expand Enabled

This optional parameter specifies whether or not a FlexGroup SnapMirror relationship and its destination FlexGroup should be auto-expanded if the source FlexGroup is expanded. This parameter is supported only for FlexGroup SnapMirror relationships. If you do not specify the parameter, the current value of auto expand will be retained. Application Consistency Group destinations are always automatically expanded when the source Application Consistency Group is expanded.

[-backoff-level {high|medium|none}] - SM Backoff Level due to Client Ops

This optional parameter specifies the SnapMirror backoff level due to client ops. This parameter is supported only for FlexVol SnapMirror relationships. The default value is `high` .

[-w, -foreground <true>] - Foreground Process

This specifies whether the operation runs as a foreground process. If this parameter is specified, the default setting is *true* (the operation runs in the foreground). When set to *true* , the command will not return until the process completes. This parameter is only applicable to relationships with *"Relationship Capability"* of *"Pre 8.2"* .

Examples

To change the schedule to *halfhour* for the SnapMirror relationship with the destination endpoint *vs2.example.com:dept_eng_dp_mirror2* , type the following command:

```
vs2.example.com::> snapmirror modify -destination-path
       vs2.example.com:dept_eng_dp_mirror2 -schedule halfhour
```

For relationships with *"Relationship Capability"* of *"Pre 8.2"* , to change the schedule to *halfhour* for the SnapMirror relationship with the destination endpoint *cluster2://vs2.example.com/dept_eng_dp_mirror2* , type the following command:

```
cluster2::> snapmirror modify -destination-path
         cluster2://vs2.example.com/dept_eng_dp_mirror2
         -schedule halfhour
```

To change the schedule to *halfhour* for the Vserver SnapMirror relationship with destination endpoint *dvs1.example.com:* , type the following command:

```
cluster2::> snapmirror modify -destination-path
         dvs1.example.com: -schedule halfhour
```

To change the policy associated with the SnapMirror Synchronous Consistency Group relationship with the destination Consistency Group *cg_dst* in Vserver *vs2.example.com* to the policy *Sync2* , type the following command:

```
vs2.example.com::> snapmirror modify -destination-path
         vs2.example.com:/cg/cg_dst -policy Sync2
```

Related Links

- snapmirror show
- snapmirror update
- snapmirror update-ls-set
- job schedule cron create
- snapmirror policy create

snapmirror promote

Promote the destination to read-write

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror promote` command performs a failover to the destination volume of a SnapMirror relationship. This command changes the destination volume from a read-only volume to a read-write volume and makes the destination volume assume the identity of the source volume. The command then destroys the original source volume. The destination volume must be a load-sharing volume. Note that you can promote a load-sharing volume that has been left in read-write state by a previously failed promote operation.

Client accesses are redirected from the original source volume to the promoted destination volume. The view clients see on the promoted destination volume is the latest transferred snapshot, which might lag behind the view clients had of the original source volume before the promote.

The SnapMirror relationship is always deleted as part of the promotion process.

It is possible that the original source volume is the source of multiple SnapMirror relationships. For such a configuration, the promoted destination volume becomes the new source volume of the other SnapMirror relationships.

This command is only supported for SnapMirror relationships with the field *"Relationship Capability"* showing as *"Pre 8.2"* in the output of the snapmirror show command.

The `snapmirror promote` command fails if a SnapMirror transfer is in progress for any SnapMirror relationship with *"Relationship Capability"* of *"Pre 8.2"* involving the original source volume. It does not fail if a SnapMirror transfer is in progress for a relationship with *"Relationship Capability"* of *"8.2 and above"*.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name*.

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name

of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name*.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

-destination-vserver <vserver name> - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

-destination-volume <volume name> - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

[-f, -force <true>] - Force

If this parameter is specified, the command proceeds without prompting for confirmation.

Examples

To promote a mirror named *vs1_root_ls1* to be the source read-write volume for mirroring and client access, type the following command:

```
cluster1::> snapmirror promote -destination-path
         //vs1.example.com/vs1_root_ls1
         -source-path //vs1.example.com/vs1_root -f true
```

Related Links

- snapmirror show

snapmirror protect

Start protection for Vservers and volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror protect` command establishes SnapMirror protection for the specified Vserver or a list of volumes. For each endpoint, the command creates a data protection destination endpoint in the Vserver specified by the `-destination-vserver` parameter, creates an extended data protection (XDP) SnapMirror relationship, and starts the initialization of the SnapMirror relationship. This command must be used from the destination Vserver or cluster. This command is not supported for FlexGroup volume constituents or non ONTAP endpoints. When source volumes are specified they must be read-write (RW) volumes.

Parameters

`[-source-cluster <Cluster name>]` - Source Cluster

This optional parameter specifies the source cluster name. This parameter is valid only when only a single Vserver is specified in the path-list parameter.

{ `-path-list {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}` - Path List

This parameter specifies the list of source endpoints to be protected. The list is a comma separated list of paths of the form *`vserver:volume`* or *`vserver:`* , for example *`vs1.example.com:dept_eng1, vs1.example.com:dept_eng2`* or *`vs1.example.com:`* . If the list contains a Vserver endpoint, then the Vserver is the only endpoint that can be specified, and the list cannot contain a mixture of volume and Vserver endpoints.

| `-source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}` - Source Path }

This optional parameter specifies the path of the source endpoint to be protected. This parameter is valid only when `-path-list` parameter is not specified. This is the alias of `-path-list` parameter.

`[-destination-vserver <vserver name>]` - Destination Vserver

This parameter specifies the Vserver in which to create the destination volumes of the SnapMirror relationships. When protecting a single Vserver, this parameter specifies the destination Vserver endpoint for protection.

`[-destination-vserver-ipspace <IPspace>]` - Destination Vserver IPSpace Name

This optional parameter specifies the IPspace the Vserver will be assigned to. If left unspecified, the Vserver will be assigned to the default IPspace. This parameter is supported while protecting a Vserver.

`[-schedule <text>]` - SnapMirror Schedule

This optional parameter designates the name of the schedule which is used to update the SnapMirror relationships.

`[-policy <sm_policy>]` - SnapMirror Policy

This parameter designates the name of the SnapMirror policy which is associated with the SnapMirror relationships. The default value is MirrorAndVault.

`[-auto-initialize {true|false}]` - Auto Initialize

This optional parameter specifies whether or not initializes of the SnapMirror relationships should be started after the relationships are created. The default value for this parameter is *`true`* .

[-destination-volume-prefix <text>] - Destination Volume Name Prefix

This optional parameter designates the prefix for the destination volume name. For example if the source path is of the form *vserver:volume* and the destination-volume-prefix specified is *prefix and no destination-volume-suffix is specified, then the destination volume name will be _prefix_volume_dst* or possibly *prefix_volume_1_dst* if a name conflict is encountered. If both prefix and suffix are specified as *prefix* and __suffix , then the destination volume name will be *prefix_volume_suffix* or *prefix_volume_1_suffix* , if a name conflict is encountered. This parameter is not supported for Vserver endpoints.

[-destination-volume-suffix <text>] - Destination Volume Name Suffix

This optional parameter designates the suffix for the destination volume name. If you do not designate a suffix, a destination volume name with suffix *dst_* will be used. For example if the source path is of the form *vserver:volume* , and the suffix specified is DP_ , the destination volume will be created with the name *volume_DP* or *volume_1_DP* if a name conflict is encountered. If both prefix and suffix are specified as *prefix and suffix* , then the destination volume name will be *prefix_volume_suffix* or *prefix_volume_1_suffix* , if a name conflict is encountered. This parameter is not supported for Vserver endpoints.

[-support-tiering {true|false}] - Provision Destination Volumes on FabricPools

This optional parameter specifies whether or not FabricPools are selected when provisioning a FlexVol volume or a FlexGroup volume during protection workflows. When this parameter is set to true, only FabricPools are used; when set to false, only non-FabricPools are used. Tiering support for a FlexVol volume can be changed by moving the volume to the required aggregate. Tiering support for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. The default value is *false* . This parameter is supported only for FlexVol volumes and FlexGroup volumes.

[-tiering-policy <Tiering Policy>] - Destination Volume Tiering Policy

This optional parameter specifies the tiering policy to apply to the destination FlexVol volume or FlexGroup volume. This policy determines whether the user data blocks of a FlexVol volume or FlexGroup volume in a FabricPool will be tiered to the capacity tier when they become cold. FabricPool combines flash (performance tier) with an object store (external capacity tier) into a single aggregate. The default tiering policy is 'snapshot-only' for a FlexVol volume and 'none' for a FlexGroup volume. The temperature of a FlexVol volume or FlexGroup volume block increases if it is accessed frequently and decreases when it is not.

The available tiering policies are:

- snapshot-only - This policy allows tiering of only the FlexVol volume or FlexGroup volume snapshots not associated with the active file system. The default minimum cooling period is 2 days. The `-tiering -minimum-cooling-days` parameter can be used to override the default using the `volume modify` command after the destination FlexVol volume or FlexGroup volume has been created.
- auto - This policy allows tiering of both snapshot and active file system user data to the capacity tier. The default cooling period is 31 days. The `-tiering-minimum-cooling-days` parameter can be used to override the default using the `volume modify` command after the destination FlexVol volume or FlexGroup volume has been created.
- none - FlexVol volume or FlexGroup volume blocks will not be tiered to the capacity tier.
- backup - On a DP FlexVol volume or FlexGroup volume this policy allows all transferred user data blocks to start in the capacity tier.

This parameter is supported only for FlexVol volumes and FlexGroup volumes.

[-destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Destination Path

This optional parameter specifies the path of the destination endpoint to be created. This parameter is valid only when a single path is specified in the `-path-list` parameter. If this parameter is specified, the `-destination-vserver` , `-destination-object-store` , `-destination-volume-prefix` , and `-destination-volume-suffix` parameters cannot be specified.

[-destination-snaplock-type {non-snaplock|compliance|enterprise}] - Destination Volume SnapLock Type

Optional parameter to create the destination endpoint with the specified SnapLock type for establishing a SnapMirror relationship. This parameter is only supported for relationships with the SnapMirror policy of type vault.

[-default-retention-period {{<integer> seconds|minutes|hours|days|months|years} | min | max | infinite | unspecified}] - Default Retention Period

Specifies the default retention period that is applied to snapshots while committing them to the locked state without an associated retention period. This parameter is only supported when the `-destination -snaplock-type` type is specified.

[-minimum-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}] - Minimum Retention Period

Specifies the minimum allowed retention period for snapshots committed to the locked state on the volume. This parameter is only supported when the `-destination-snaplock-type` type is specified.

[-maximum-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}] - Maximum Retention Period

Specifies the maximum allowed retention period for snapshots committed to the locked state on the volume. This parameter is only supported when the `-destination-snaplock-type` type is specified.

Examples

To establish SnapMirror protection for the source volumes *`vs1.example.com:dept_eng1`* and *`vs1.example.com:dept_eng2`* using destination-vserver *`vs2.example.com`* and policy *`MirrorAllSnapshots`* type the following command:

```
vs2.example.com::> snapmirror protect -path-list
       vs1.example.com:dept_eng1,vs1.example.com:dept_eng2 -destination
-vserver
       vs2.example.com -policy MirrorAllSnapshots
```

To establish SnapMirror protection for the source Vserver *`vs1.example.com`* which is on cluster *`cluster1`* creating a destination-vserver named *`vs1dp.example.com`* and using policy *`MirrorAllSnapshots`* type the following command:

```
cluster2::> snapmirror protect -source-cluster cluster1 -path-list
vs1.example.com: -destination-vserver vs1dp.example.com -policy
MirrorAllSnapshots
```

To establish SnapMirror protection for the source Consistency Group endpoint *vs1.example.com:/cg/cg1* which is on cluster *cluster1* creating a destination Consistency Group *vs2.example.com:/cg/cg1_dst* and using policy *MirrorAndVault* type the following command:

```
cluster2::> snapmirror protect -path-list vs1.example.com:/cg/cg1 -policy
MirrorAndVault
```

snapmirror quiesce

Disable future transfers

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror quiesce` command disables future transfers for a SnapMirror relationship. If there is no transfer in progress, the relationship becomes *"Quiesced"* .

If there is a transfer in progress, it is not affected, and the relationship becomes *"Quiescing"* until the transfer completes. If the current transfer aborts, it will be treated like a future transfer and will not restart.

If applied to a load-sharing (LS) SnapMirror relationship, all the relationships in the load-sharing set will be quiesced.

If applied to a relationship with a policy of type *strict-sync-mirror* or *sync-mirror* , any ongoing operation is aborted, and if the relationship is InSync, synchronous replication is stopped. If the replication policy type is "strict-sync-mirror", then the primary client I/O disruption is not enforced. A new common snapshot is created if the relationship is InSync, unless a recent one exists. The relationship becomes *"Quiescing"* until these operations complete.

When a SnapMirror relationship is quiesced, it remains quiesced across reboots and fail-overs.

This command is supported for SnapMirror relationships with the field *"Relationship Capability"* showing as either *"8.2 and above"* or *"Pre 8.2"* in the output of the snapmirror show command.

This command is not supported for SnapMirror active sync relationships with policy of type *automated-failover* or *automated-failover-duplex* .

The `snapmirror quiesce` command must be used from the destination Vserver or cluster.

The relationship must exist on the destination Vserver or cluster. When issuing `snapmirror quiesce` , you must specify the destination endpoint. The specification of the source endpoint of the relationship is optional.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

This parameter specifies the source endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault source, the source endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire source, the source endpoint is specified in the form *hostip:/lun/name*.

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the source cluster of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-source-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameters `-source-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-source-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

This parameter specifies the destination endpoint of the SnapMirror relationship in one of four path formats. The normal format includes the names of the Vserver (vserver) and/or volume (volume). To support relationships with *"Relationship Capability"* of *"Pre 8.2"*, a format which also includes the name of the cluster (cluster) is provided. The *"Pre 8.2"* format cannot be used when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*. For SnapMirror relationships with an AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For SnapMirror relationships with a SolidFire destinations, the destination endpoint is specified in the form *hostip:/lun/name*.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the destination cluster of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and `-destination-volume` must also be specified. This parameter is only applicable for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter cannot be specified when operating in a Vserver context on relationships with *"Relationship Capability"* of *"8.2 and above"*.

`-destination-vserver <vserver name>` - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameters `-destination-volume` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

`-destination-volume <volume name>` - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameters `-destination-vserver` and for relationships with *"Relationship Capability"* of *"Pre 8.2"*, `-destination-cluster` must also be specified. This parameter is not supported for relationships with non-ONTAP destination endpoints.

Examples

To quiesce the SnapMirror relationship with the destination endpoint *vs2.example.com:dept_eng_mirror2*, type the following command:

```
vs2.example.com::> snapmirror quiesce -destination-path
       vs2.example.com:dept_eng_mirror2
```

For relationships with *"Relationship Capability"* of *"Pre 8.2"*, to quiesce the SnapMirror relationship with the destination endpoint *cluster2://vs2.example.com/dept_eng_mirror2*, type the following command:

```
cluster2::> snapmirror quiesce -destination-path
         cluster2://vs2.example.com/dept_eng_mirror2
```

To quiesce the Vserver SnapMirror relationship with the destination endpoint *dvs1.example.com:*, type the following command:

```
cluster2::> snapmirror quiesce -destination-path
         dvs1.example.com:
```

To quiesce the Application Consistency Group SnapMirror relationship with the destination Application Consistency Group *app_cg_dst* in Vserver *vs2.example.com*, type the following command:

```
vs2.example.com::> snapmirror quiesce -destination-path
          vs2.example.com:/cg/app_cg_dst
```

Related Links

- snapmirror show

snapmirror release

Remove source information for a SnapMirror relationship

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror release` command removes the relationship information from the source Vserver. The command also removes any snapshot owner tags and any snapshots which were created for the specified relationship from the source volumes. It does not destroy any volumes or Vservers. This command must be used from the source Vserver or cluster.

For relationships with a policy of type *`strict-sync-mirror`* , it additionally resumes primary volume access if the IO was restricted because the relationship was OutOfSync.

For relationship having Application Consistency Groups, `snapmirror release` command used on the source of a Application Consistency Group relationship, removes the relationship information from the source side and deletes snapshots which were created on the source volumes.

You can use the snapmirror list-destinations command to display source Vservers' relationship information. This information is populated during the first SnapMirror transfer, not when the snapmirror create command is issued.

This command is not supported for SnapMirror relationships with the field *`"Relationship Capability"`* showing as *`"Pre 8.2"`* in the output of the snapmirror show command.

This command is not supported for SnapMirror relationships with non-ONTAP endpoints.

The `snapmirror release` operation fails if a SnapMirror transfer for the SnapMirror relationship is in a data phase of the transfer.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

Specifies the source endpoint of the SnapMirror relationship in one of three formats. The normal format includes the names of the Vserver (vserver), and/or volume (volume). A format which also includes the name of the cluster (cluster) is also provided for consistency with other SnapMirror commands. The form of the pathname which includes the cluster name cannot be used when operating in a Vserver context. For relationships between Application Consistency Groups, the source endpoint is specified in the form *`[vserver:]/cg/cg-name`* .

| [-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameter `-source-volume` must also be specified.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, parameter `-source-vserver` must also be specified.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

Specifies the destination endpoint of the SnapMirror relationship in one of three formats. The normal format includes the names of the Vserver (vserver), and/or volume (volume). A format which also includes the name of the cluster (cluster) is also provided for consistency with other SnapMirror commands. The form of the pathname which includes the cluster name cannot be used when operating in a Vserver context. For relationships between Application Consistency Groups the destination endpoint is specified in the form *[vserver:]/cg/cg-name*

| -destination-vserver <vserver name> - Destination Vserver

Specifies the destination Vserver of the SnapMirror relationship. For relationships with volumes as endpoints, if this parameter is specified, parameter `-destination-volume` must also be specified.

[-destination-volume <volume name>] - Destination Volume }

Specifies the destination volume of the SnapMirror relationship. If this parameter is specified, parameter `-destination-vserver` must also be specified.

[-relationship-info-only <true>] - Remove relationship info only (skip cleanup of snapshots)

If this parameter is specified, the cleanup of snapshots is bypassed and only the source relationship information is removed. It is recommended to specify this parameter only when the source volume is not accessible.

[-relationship-id <UUID>] - Relationship ID

This optional parameter specifies the relationship identifier of the relationship. It must be specified when information for more than one relationship with the same source and destination paths is present. This parameter is not supported for Vserver SnapMirror relationships.

[-f, -force <true>] - Force

If this parameter is specified, the command proceeds without prompting for confirmation.

Examples

To release the source information for the SnapMirror relationship with the destination endpoint *vs2.example.com:dept_eng_dp_mirror4* , type the following command:

```
vs1.example.com::> snapmirror release
        -destination-path vs2.example.com:dept_eng_dp_mirror4
```

To release the source information for the SnapMirror relationship with the destination endpoint *vs2.example.com:dept_eng_dp_mirror4* , and relationship-id *5f91a075-6a72-11e1-b562-123478563412* , type the following command:

```
vs1.example.com::> snapmirror release
       -destination-path vs2.example.com:dept_eng_dp_mirror4
       -relationship-id 5f91a075-6a72-11e1-b562-123478563412
```

To release the source information for the SnapMirror relationship with the destination endpoint *dvs1.example.com:* , type the following command:

```
cluster1::> snapmirror release
       -destination-path dvs1.example.com:
```

To release the Application Consistency Group SnapMirror relationship with the destination Application Consistency Group *cg_dst* in Vserver *vs2.example.com* and the source Application Consistency Group *cg_src* in Vserver *vs1.example.com* , type the following command from the source cluster:

```
source::> snapmirror release
        vs2.example.com:/cg/cg_dst
```

To release just the source information but not remove the snapshots that might be needed for a subsequent resync for the Asynchronous SnapMirror Application Consistency Group relationship with the destination Application Consistency Group *cg_dst* in Vserver *vs2.example.com* , type the following command:

```
vs2.example.com::> snapmirror release
       -destination-path vs2.example.com:/cg/cg_dst
       -relationship-info-only true
```

To release the SnapMirror active sync relationship with the destination Consistency Group *cg_dst* in Vserver *vs2.example.com* and the source Consistency Group *cg_src* in Vserver *vs1.example.com* , type the following command from the source cluster:

```
source::> snapmirror release
        vs2.example.com:/cg/cg_dst
```

Related Links

- snapmirror list-destinations
- snapmirror create
- snapmirror show

snapmirror restore

Restore a Snapshot copy from a source volume to a destination volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `snapmirror restore` command restores the entire contents of a snapshot or one or more files, LUNs or NVMe namespaces of a snapshot from one volume to another volume.

The `snapmirror restore` command, in context to Application Consistency Groups, restores the entire content of a snapshot from one Application Consistency Group to another Application Consistency Group. Each Application Consistency Group is considered as atomic entity in restore operation

The source of the restore can be a volume or an Application Consistency Group that is:

- the destination volume or destination Application Consistency Group of a extended data protection (XDP) relationship
- the destination volume of a data protection (DP) relationship with *"Relationship Capability"* of *"8.2 and above"*
- a data-protection volume or an Application Consistency Group which is not the destination endpoint of any SnapMirror relationship
- a read-write volume or an Application Consistency Group.
- an AltaVault endpoint. In this case the destination must be an empty ONTAP volume.

The following cannot be used as either the source or destination volume or destination Application Consistency Group of a restore:

- a volume or Application Consistency Group that is the source or destination endpoint of a SnapMirror load-sharing relationship.
- a volume that is the destination endpoint of a SnapMirror relationship with the *"Relationship Capability"* of *"Pre 8.2"*.
- a SolidFire endpoint.

A SnapMirror relationship of type *RST* is created from the source volume or source Application Consistency Group to the destination volume or destination Application Consistency Group by the `snapmirror restore` command. This relationship lasts for the duration of the restore operation and is deleted when the command completes successfully.

The following paragraphs describe the behavior when restoring the entire contents of a snapshot to a destination volume or destination Application Consistency Group.

By default the `snapmirror restore` will copy the latest snapshot from the source volume or source Application Consistency Group to the destination volume or destination Application Consistency Group. A specific snapshot can be selected with the `-source-snapshot` parameter.

Any quota rules defined for the destination volume are deactivated prior to restoring the entire contents of a snapshot. Run the command volume quota modify`-vserver` *destination-volume-vserver*-volume *destination-volume-name*-state *on* to reactivate quota rules after the entire contents of the snapshot have been restored.

If the destination Application Consistency Group contains empty data protection volumes, the `snapmirror restore` command performs a baseline restore.

If the destination volume is an empty data protection volume, the `snapmirror restore` command performs a baseline restore. For a baseline restore the following steps are performed:

- Create the *RST* SnapMirror relationship.
- The entire contents of the snapshot selected to be restored are copied to the active file system of the destination volume or destination Application Consistency Group.

- The destination volume or destination Application Consistency Group is made read-write.
- The *RST* SnapMirror relationship is deleted.

If the destination volume is a read-write volume, an incremental restore is performed. The incremental restore fails if it cannot find a common snapshot between the source and destination volumes. Restoring a snapshot to an empty read-write volume is not supported. Incremental restore from a non-ONTAP endpoint to a ONTAP volume is not supported.

If the destination Application Consistency Group contains a read-write flex volume, an incremental restore is performed. The incremental restore fails if it cannot find a common snapshot between the source and destination Application Consistency Group. Restoring a snapshot to an empty read-write flex volumes which are the part of Application Consistency Group is not supported.

An incremental restore preserves all snapshots on the destination volume but does not preserve changes to the active file system since the latest snapshot. To preserve changes to the destination volume since the latest snapshot use the volume snapshot create command. Restore is a disruptive operation so client access of the destination volume is not advised for the duration of the operation.

For an incremental restore the following steps are performed:

- Create the *RST* SnapMirror relationship.
- The active file system of the destination volume is reverted to the latest snapshot on the destination volume and the destination volume is made read-only.
- In context to Application Consistency Group, the active file system of the destination Application Consistency Group is reverted to the latest snapshot on the destination Application Consistency Group and the flex volumes inside destination Application Consistency Group is made read-only.
- This snapshot is the exported snapshot and it is the view to which clients are redirected when accessing the destination volume or destination Application Consistency Group.
- The contents of the snapshot selected to be restored are copied to the active file system of the destination volume or destination Application Consistency Group.
- The destination volume or destination Application Consistency Group is made read-write.
- The *RST* SnapMirror relationship is deleted.

If `snapmirror restore` fails or is aborted, the *RST* relationship remains. Use the snapmirror show command with the destination volume name or destination Application Consistency Group name to display the reason for the error. An EMS is also generated when a failure occurs. There are two options to recover when restore fails or is aborted:

- Take corrective action suggested by the EMS and reissue the original command.
- Use the original command with `-clean-up-failure` to cancel the request.

When specifying `-clean-up-failure` to cancel an incremental restore request, the following steps are performed:

- If the snapshot has not been restored to the destination volume or destination Application Consistency Group, all data copied to the active file system by `snapmirror restore` to the destination volume or destination Application Consistency Group is reverted.
- The destination volume or destination Application Consistency Group is made read-write.
- The *RST* SnapMirror relationship is deleted.

When specifying `-clean-up-failure` to cancel a baseline restore request, the following steps are performed:

- If the snapshot has been restored to the destination volume or destination Application Consistency Group, the volume or Application Consistency Group is made read-write.
- The *`RST`* SnapMirror relationship is deleted.

The following paragraphs describe the behavior and requirements when restoring one or more files, LUNs or NVMe namespaces to the destination volume.

The destination volume must be a read-write volume. Restoring files, LUNs or NVMe namespaces to a data protection volume is not supported. When restoring files, LUNs or NVMe namespaces the source and destination volumes are not required to have a common snapshot. If a common snapshot exists, an incremental restore is performed for those files, LUNs or NVMe namespaces being restored which exist in the common snapshot.

The destination Application Consistency Group must contains a read-write volume. While restoring, source and destination Application Consistency Group are not required to have a common snapshot. If a common snapshot exists, an incremental restore is performed.

The contents of the files, LUNs or NVMe namespaces to which data is being restored on the destination volume are not preserved by this command. To preserve the contents of the destination files, LUNs or NVMe namespaces, create a snapshot on the destination volume prior to running this command. Client I/O is not allowed to a file, LUN or NVMe namespace to which data is being restored on the destination volume.

The `-source-snapshot` parameter is required when restoring files, LUNs or NVMe namespaces. It identifies the snapshot on the source volume from which the files, LUNs or NVMe namespaces to be restored are copied. If all files, LUNs or NVMe namespaces to be restored do not exist in this snapshot the command fails.

The source path for each file, LUN or NVMe namespace being restored is required. The source path of a file, LUN or NVMe namespace is from the root of the source snapshot of the source volume. The file is restored to the same path on the destination volume unless an optional destination path is specified. The destination path is from the root of the destination volume. If a file, LUN or NVMe namespace to which data is being restored on the destination volume does not exist, the file, LUN or NVMe namespace is created. If any directory in the path of the file, LUN or NVMe namespace being restored does not exist on the destination volume, the command fails. Overwriting the contents of an existing file with the contents of a different file is supported. Similarly, overwriting the contents of an existing LUN or NVMe namespace with the contents of a different LUN or NVMe namespace is supported. However, overwriting a file with the contents of a LUN or NVMe namespace is not supported. Overwriting a LUN with the contents of a file or NVMe namespace is not supported. Overwriting an NVMe namespace with the contents of a file or LUN is not supported. Client I/O is not allowed to all files, LUNs and NVMe namespaces to which data is being restored on the destination volume.

If quota rules have been defined for the destination volume, resource usage is updated during file restore, but limits of quota rules are not enforced. Therefore, resource limits might be exceeded during a file restore.

Multiple concurrent `snapmirror restore` commands, restoring one or more files, LUNs or NVMe namespaces to the same destination volume, are not supported. The destination volume of a `snapmirror restore` to which one or more files, LUNs or NVMe namespaces are being restored, can simultaneously be the source volume of a snapmirror update .

For a file, LUN or NVMe namespace restore the following steps are performed:

- Create the *`RST`* SnapMirror relationship.

- If any file, LUN or NVMe namespace being restored does not exist on the destination volume, create all such files, LUNs or NVMe namespaces.
- Prevent client I/O to files, LUNs or NVMe namespaces to which data is being restored on the destination volume.
- Revoke locks and space reservations held by NAS clients for files being restored.
- Copy the contents of all source files, LUNs or NVMe namespaces to the corresponding file, LUN or NVMe namespace on the destination volume.
- Allow client I/O to files, LUNs or NVMe namespaces to which data has been restored on the destination volume.
- Delete the *RST* SnapMirror relationship.

> Some file restore operations require a snapshot to be created. This snapshot is temporary, it is deleted before the operation completes.

Since client I/O is not allowed to files, LUNs or NVMe namespaces being restored, client I/O to files, LUNs or NVMe namespaces being restored should be quiesced. Mapped LUNs or NVMe namespaces remain mapped throughout the operation. SAN clients do not need to rediscover a mapped LUN that has been restored. Restoring an NVMe namespace on top of another NVMe namespace with a different attribute relevant to NVMe protocol accessibility (like size) is not supported.

If `snapmirror restore` fails or is aborted, the *RST* relationship remains. Use the snapmirror show command with the destination volume or destination Application Consistency Group to display the reason for the error. An EMS is also generated when a failure occurs. There are two options to recover when restore fails or is aborted:

- Take corrective action suggested by the EMS and reissue the original command.
- Use `snapmirror restore`-clean-up-failure` along with specifying the destination volume or destination Application Consistency Group to cancel the request.

When specifying `-clean-up-failure` to cancel a file restore request, the following steps are performed:

- Any files to which client I/O is not allowed are removed.
- Any snapshot created for use during a file restore operation is deleted.
- The *RST* SnapMirror relationship is deleted.

> LUNs to which client I/O is not allowed remain. For LUNs to which client I/O is not allowed, do one of the following:

- Use the `snapmirror restore` command to restore data to the LUN. Once the command completes successfully, client I/O to the LUN is allowed.
- Delete the LUN using the lun delete command with the `-force-fenced` parameter.

> Similarly, NVMe namespaces to which client I/O is not allowed remain. For NVMe namespaces to which client I/O is not allowed, do one of the following:

- Use the `snapmirror restore` command to restore data to the NVMe namespace. Once the command completes successfully, client I/O to the NVMe namespace is allowed.
- Delete the NVMe namespace using the vserver nvme namespace delete command with the `-skip`

`-mapped-check` parameter.

The `snapmirror restore` command must be used from the destination Vserver or cluster.

With `-lun-list` , subset of LUNs can be restored to the Destination Consistency Group endpoint.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}] - Source Path

Specifies the source endpoint in one of four formats. The basic format includes the names of the Vserver (vserver) and volume (volume). A format which also includes the name of the cluster (cluster) is supported for consistency with other `snapmirror` commands. The form of the pathname which includes the cluster name is not valid when operating in a Vserver context. Source endpoint having Application Consistency Group includes vserver name followed by *`cg`* and Application Consistency Group name. A non-ONTAP source endpoint (for example, AltaVault) can be specified in the form *`hostip:/share/share-name`* .

| [-source-cluster <Cluster name>] - Source Cluster

Specifies the cluster in which the source volume resides. This parameter is not needed; it is provided for consistency with other `snapmirror` commands. If this parameter is specified, the `-source-vserver` and `-source-volume` parameters must also be specified. This parameter is not valid when operating in a Vserver context. This parameter is not supported if the source is a non-ONTAP endpoint.

[-source-vserver <vserver name>] - Source Vserver

Specifies the source Vserver of the SnapMirror relationship. If this parameter is specified, the `-source-volume` parameter must also be specified. This parameter is not supported if the source is a non-ONTAP endpoint.

[-source-volume <volume name>] - Source Volume }

Specifies the source volume of the SnapMirror relationship. If this parameter is specified, the `-source-vserver` parameter must also be specified. This parameter is not supported if the source is a non-ONTAP endpoint.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>} - Destination Path

Specifies the destination endpoint in one of three formats. The basic format includes the names of the Vserver (vserver) and volume (volume). A format that also includes the name of the cluster (cluster) is supported for consistency with other `snapmirror` commands. The form of the pathname which includes the cluster name is not valid when operating in a Vserver context. Destination endpoint having Application Consistency Group includes vserver name followed by *`cg`* and Application Consistency Group name.

| [-destination-cluster <Cluster name>] - Destination Cluster

Specifies the cluster in which the destination volume resides. This parameter is not needed; it is provided for consistency with other `snapmirror` commands. If this parameter is specified, the `-destination-vserver` and `-destination-volume` parameters must also be specified. This parameter is not valid when operating in a Vserver context. This parameter is only applicable for relationships with *`"Relationship Capability"`* of *`"Pre 8.2"`*.

-destination-vserver <vserver name> - Destination Vserver

Specifies the destination Vserver. If this parameter is specified, the `-destination-volume` parameter must also be specified.

-destination-volume <volume name> - Destination Volume }

Specifies the destination volume. If this parameter is specified, the `-destination-vserver` parameter must also be specified.

[-s, -source-snapshot <text>] - Source Snapshot

When restoring the entire contents of a snapshot, this optional parameter identifies the snapshot to be restored from the source volume to the destination volume. The default value is the latest snapshot on the source volume. When restoring one or more files, LUNs or NVMe namespaces from a snapshot, this parameter is required.

[-k, -throttle <throttleType>] - Throttle (KB/sec)

This optional parameter limits the network bandwidth used for the restore transfer when the source and destination volumes belong to different clusters. It sets the maximum rate (in Kbytes/sec) at which data can be transferred between the clusters during the operation. To fully use the network bandwidth available between the clusters, set the throttle value to *`unlimited`* or *`0`* . The minimum throttle value is four Kbytes/sec, so if you specify a throttle value between *`1`* and *`4`* , it will be treated as if you specified *`4`* .

[-transfer-priority {low|normal}] - Transfer Priority

This optional parameter specifies the priority at which the transfer runs. The default value for this parameter is `normal` .

[-cg-item-mappings {<source volume>:@<destination volume>|<source item path>:@<destination item path>}] - Consistency Group Item Mappings

This optional parameter specifies a list of the consistency group (CG) item mappings. It is considered only if the supplied source and destination path values denote valid CG paths. For CG level relationships, this parameter must be specified. The value for this parameter must contain mappings of one or more pairs of constituent source and destination volumes of the form *`srcvol1:@dstvol1,…`* .

[-disable-storage-efficiency <true>] - Disable Storage Efficient Transfer

The default behavior of restore is to preserve storage efficiency when possible. Use this optional parameter to turn off storage efficiency for data transferred over the wire and written to the destination volume.

[-clean-up-failure <true>] - Clean up after Failure

Use this optional parameter to recover from an aborted or failed restore operation. Any temporary *`RST`* relationship is removed from the destination Vserver. An attempt is made to remove any temporary *`RST`* relationship from the source Vserver. If cleaning up an incomplete restore of the entire contents of a snapshot and the destination volume was read-write prior to the failed or aborted restore operation, it is converted back to read-write if necessary, while removing all data transferred or copied during the restore operation. If cleaning up an incomplete restore of one or more files, LUNs or NVMe namespaces of a snapshot, any file to which client I/O is not allowed is deleted.

[-tries <integer_or_unlimited>] - Tries Limit

Specifies the total number of attempts to transfer data in cases where a transfer is interrupted by an error that SnapMirror can recover from. The value of this parameter must be a positive integer or *`unlimited`* .

[-f, -force <true>] - Force

If this parameter is specified, the command proceeds without prompting for confirmation.

{ [-file-list <<source path>[,@<destination path>]>,…] - File List

Specifies the files, LUNs or NVMe namespaces to be restored. The list can contain specifications for up to 8 files, LUNs or NVMe namespaces. Specification for each file, LUN or NVMe namespace consists of a *source_path* and an optional *destination_path* , and is of the form '*source_path[,@destination_path*]'. *source_path* is required and is the path of the file from the source snapshot, e.g. /dira/file1 or /lun1. The source path does not include the snapshot name nor the source volume name. The path to each file to be restored in the active file system of the destination volume is the same as the path specified by *source_path* , unless an optional *destination_path* is specified. *destination_path* begins with the @ symbol followed by the path of the file from the root of the active file system of the destination volume, e.g. @/file1 or @/dira/lun1. File restore is not supported if the *source_path* or *destination_path* contains commas in its directory or file name. Each *source_path* and *destination_path* is a separate entity in the list of paths. A *destination_path* is associated with the *source_path* that immediately precedes it. If this parameter is specified, `-source-snapshot` must also be specified. Examples:

```
/dira/file1
```

```
/dira/file1,@/dirb/file2
```

```
/dira/file1,@/dirb/file2,/dirc/file3
```

| [-lun-list [source lun>:@<destination lun],…] - Lun List }

Specifies the source and destination LUN pairs to be restored. Specification for each LUN list consists of a *source_lun* and *destination_lun* , and is of the form '*source_lun:@destination_lun* '. *source_lun* is required and is the name of the LUN from the source snapshot. A source LUN name is considered as destination LUN if destination LUNs are not provided. Examples:

```
source_lun1
```

:@

```
destination_lun1
```

,or

source_lun1

```
,
```

source_lun2

```
``[-use-network-compression <true>]`` - Use Network Compression::
Use this optional parameter to use network compression for data transfer
over the wire. This parameter is not supported for relationships with non-
ONTAP endpoints.

``[-backoff-level {high|medium|none}]`` - SM Backoff Level due to Client
Ops::
This optional parameter specifies the SnapMirror backoff level due to
client ops. This parameter is supported only for FlexVol SnapMirror
relationships. The default value is `high` .

== Examples

The following example does an incremental restore between the restore
source volume ``_vs2.example.com:dept_eng_dp_mirror2_`` and the restore
destination volume ``_vs1.example.com:dept_eng_`` :
```

vs1.example.com::> snapmirror restore
-destination-path vs1.example.com:dept_eng
-source-path vs2.example.com:dept_eng_dp_mirror2
-source-snapshot snap3
Warning: All data newer than snapshot snap6 on volume
vs1.example.com:dept_eng will be deleted.
Do you want to continue? {y|n}: y
[Job 34] Job is queued: snapmirror restore from source
vs2.example.com:dept_eng_dp_mirror2 for the snapshot snap3.
vs1.example.com::>

```
The following example does a restore between the source Application
Consistency Group ``_vs2.example.com:/cg/src_cg_`` and destination
Application Consistency Group ``_vs1.example.com:/cg/dst_cg_`` :
```

vs1.example.com::> snapmirror restore
-destination-path vs1.example.com:/cg/dst_cg
-source-path vs2.example.com:/cg/src_cg
-source-snapshot snap3
[Job 34] Job is queued: snapmirror restore from source
vs2.example.com:/cg/src_cg for the snapshot snap3.
vs1.example.com::>

```
The following example restores /file3 from the source snapshot ``_snap3_``
on the source volume ``_vs2.example.com:dept_eng_dp_mirror2_`` to the
active file system of the restore destination volume
``_vs1.example.com:dept_eng_`` :
```

vs1.example.com::> snapmirror restore
-destination-path vs1.example.com:dept_eng
-source-path vs2.example.com:dept_eng_dp_mirror2
-source-snapshot snap3
-file-list /file3
Warning: This command will overwrite any file on destination
"vs1.example.com:dept_eng" that has the same path as any of
the files to be restored.
Do you want to continue? {y|n}: y
[Job 35] Job is queued: snapmirror restore from source
"vs2.example.com:dept_eng_dp_mirror2" for the snapshot snap3.
vs1.example.com::>

```
The following example restores /file3 from the source snapshot ``_snap3_``
on the source volume ``_vs2.example.com:dept_eng_dp_mirror2_`` to
/file3.new in the active file system of the restore destination volume
``_vs1.example.com:dept_eng_`` :
```

vs1.example.com::> snapmirror restore
-destination-path vs1.example.com:dept_eng
-source-path vs2.example.com:dept_eng_dp_mirror2
-source-snapshot snap3
-file-list /file3,@/file3.new
Warning: This command will overwrite any file on destination
"vs1.example.com:dept_eng" that has the same path as any of
the files to be restored.
Do you want to continue? {y|n}: y
[Job 36] Job is queued: snapmirror restore from source
"vs2.example.com:dept_eng_dp_mirror2" for the snapshot snap3.
vs1.example.com::>

```
The following example restores /file1, /file2, and /file3 from the source
snapshot ``_snap3_`` on the source volume
``_vs2.example.com:dept_eng_dp_mirror2_`` respectively to /file1.new,
/file2, and /file3.new in the active file system of the restore
destination volume ``_vs1.example.com:dept_eng_`` :
```

vs1.example.com::> snapmirror restore
-destination-path vs1.example.com:dept_eng
-source-path vs2.example.com:dept_eng_dp_mirror2
-source-snapshot snap3
-file-list /file1,@/file1.new,/file2,/file3,@/file3.new

Warning: This command will overwrite any file on destination
"vs1.example.com:dept_eng" that has the same path as any of
the files to be restored.
Do you want to continue? {y|n}: y
[Job 36] Job is queued: snapmirror restore from source
"vs2.example.com:dept_eng_dp_mirror2" for the snapshot snap3.
vs1.example.com::>

```
== Related Links

* xref:{relative_path}volume-quota-modify.html[volume quota modify]
* xref:{relative_path}volume-snapshot-create.html[volume snapshot create]
* xref:{relative_path}snapmirror-show.html[snapmirror show]
* xref:{relative_path}snapmirror-update.html[snapmirror update]
* xref:{relative_path}lun-delete.html[lun delete]
* xref:{relative_path}vserver-nvme-namespace-delete.html[vserver nvme
namespace delete]




[[IDc31327491e175b90cd968ae95b6cce6d]]
= snapmirror resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable future transfers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror resume` command enables future transfers for a SnapMirror
relationship that has been quiesced.

If there is a scheduled transfer for the relationship, it will be
triggered on the next schedule. If there is a restart checkpoint, it will
be re-used if possible.

If applied on a load-sharing (LS) SnapMirror relationship, it enables
future transfers for all the relationships in the load-sharing set.

If applied on a relationship with a policy of type ``_strict-sync-
```

```
mirror_`` or ``_sync-mirror_`` , it enables future resync operations and
initiates an Auto Resync.

When a quiesced SnapMirror relationship is resumed, future transfers
remain enabled across reboots and fail-overs.

This command is supported for SnapMirror relationships with the field
``_"Relationship Capability"_`` showing as either ``_"8.2 and above"_`` or
``_"Pre 8.2"_`` in the output of the xref:{relative_path}snapmirror-
show.html[snapmirror show] command.

The `snapmirror resume` command must be used from the destination Vserver
or cluster.

The relationship must exist on the destination Vserver or cluster. When
issuing `snapmirror resume` , you must specify the destination endpoint.
The specification of the source endpoint of the relationship is optional.

== Parameters

{ ``[-S, -source-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Source Path::
This parameter specifies the source endpoint of the SnapMirror
relationship in one of four path formats. The normal format includes the
names of the Vserver (vserver) and/or volume (volume). To support
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , a
format which also includes the name of the cluster (cluster) is provided.
The ``_"Pre 8.2"_`` format cannot be used when operating in a Vserver
context on relationships with ``_"Relationship Capability"_`` of ``_"8.2
and above"_`` . For SnapMirror relationships with an AltaVault source, the
source endpoint is specified in the form ``_hostip:/share/share-name_`` .
For SnapMirror relationships with a SolidFire source, the source endpoint
is specified in the form ``_hostip:/lun/name_`` .

| ``[-source-cluster <Cluster name>]`` - Source Cluster::
Specifies the source cluster of the SnapMirror relationship. If this
parameter is specified, the `-source-vserver` and `-source-volume`
parameters must also be specified. This parameter is only applicable for
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .
This parameter cannot be specified when operating in a Vserver context on
relationships with ``_"Relationship Capability"_`` of ``_"8.2 and
above"_`` .

``[-source-vserver <vserver name>]`` - Source Vserver::
Specifies the source Vserver of the SnapMirror relationship. For
```

```
relationships with volumes as endpoints, if this parameter is specified,
parameters `-source-volume` and for relationships with ``_"Relationship
Capability"_`` of ``_"Pre 8.2"_`` , `-source-cluster` must also be
specified. This parameter is not supported for relationships with non-
ONTAP source endpoints.

``[-source-volume <volume name>]`` - Source Volume }::
Specifies the source volume of the SnapMirror relationship. If this
parameter is specified, parameters `-source-vserver` and for relationships
with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-source-
cluster` must also be specified. This parameter is not supported for
relationships with non-ONTAP source endpoints.

{ ``-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}`` - Destination
Path::
This parameter specifies the destination endpoint of the SnapMirror
relationship in one of four path formats. The normal format includes the
names of the Vserver (vserver) and/or volume (volume). To support
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , a
format which also includes the name of the cluster (cluster) is provided.
The ``_"Pre 8.2"_`` format cannot be used when operating in a Vserver
context on relationships with ``_"Relationship Capability"_`` of ``_"8.2
and above"_`` . For SnapMirror relationships with an AltaVault
destinations, the destination endpoint is specified in the form
``_hostip:/share/share-name_`` . For SnapMirror relationships with a
SolidFire destinations, the destination endpoint is specified in the form
``_hostip:/lun/name_`` .

| ``[-destination-cluster <Cluster name>]`` - Destination Cluster::
Specifies the destination cluster of the SnapMirror relationship. If this
parameter is specified, parameters `-destination-vserver` and `-
destination-volume` must also be specified. This parameter is only
applicable for relationships with ``_"Relationship Capability"_`` of
``_"Pre 8.2"_`` . This parameter cannot be specified when operating in a
Vserver context on relationships with ``_"Relationship Capability"_`` of
``_"8.2 and above"_`` .

``-destination-vserver <vserver name>`` - Destination Vserver::
Specifies the destination Vserver of the SnapMirror relationship. For
relationships with volumes as endpoints, if this parameter is specified,
parameters `-destination-volume` and for relationships with
``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-destination-
cluster` must also be specified. This parameter is not supported for
relationships with non-ONTAP destination endpoints.
```

```
``-destination-volume <volume name>`` - Destination Volume }::
Specifies the destination volume of the SnapMirror relationship. If this
parameter is specified, parameters `-destination-vserver` and for
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-
destination-cluster` must also be specified. This parameter is not
supported for relationships with non-ONTAP destination endpoints.

== Examples

To re-enable future transfers for the SnapMirror relationship with the
destination endpoint ``_vs2.example.com:dept_eng_dp_mirror2_`` that has
been previously quiesced, type the following command:
```

vs2.example.com::> snapmirror resume -destination-path vs2.example.com:dept_eng_dp_mirror2

```
To re-enable future transfers for the SnapMirror relationship with the
destination endpoint ``_cluster2://vs2.example.com/dept_eng_dp_mirror2_``
that has been previously quiesced, type the following command:
```

cluster2::> snapmirror resume -destination-path cluster2://vs2.example.com/dept_eng_dp_mirror2

```
To re-enable future transfers for the Vserver SnapMirror relationship with
the destination endpoint ``_dvs1.example.com:_`` that has been previously
quiesced, type the following command:
```

cluster2::> snapmirror resume -destination-path dvs1.example.com:

```
To re-enable future transfers of the Application Consistency Group
SnapMirror relationship with the destination Consistency Group
``_app_cg_dst_`` in Vserver ``_vs2.example.com_`` , type the following
command:
```

vs2.example.com::> snapmirror resume -destination-path vs2.example.com:/cg/app_cg_dst

```
== Related Links

* xref:{relative_path}snapmirror-show.html[snapmirror show]
```

```
[[IDabf5db8ef0ab00a71a5c1eb29303828e]]
= snapmirror resync
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a resynchronize operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror resync` command establishes or reestablishes a mirroring
relationship between a source and a destination endpoint. The endpoints
can be Vservers, volumes, Application Consistency Groups or non-ONTAP
endpoints that support SnapMirror. `snapmirror resync` for a SnapMirror
relationship with volumes as endpoints is typically executed in the
following cases:

* The destination mirror is broken (that is, the destination volume is a
read-write volume and no longer a data protection mirror). After the
`snapmirror resync` command completes, the destination volume is made a
data protection mirror and the mirror can be manually updated or scheduled
for updates.
* xref:{relative_path}snapmirror-update.html[snapmirror update] command
failed because the required common snapshot was deleted on the source
volume.
* The volumes are the first and third endpoints in a cascade chain of
relationships and they have a common snapshot. In this case, `snapmirror
resync` might implicitly create the SnapMirror relationship between them.



NOTE: The `snapmirror resync` command can cause data loss on the
destination volume because the command can remove the exported snapshot on
the destination volume.



The default behavior of the `snapmirror resync` command for volume
relationships is defined as follows:
```

```
* Finds the most recent common snapshot between the source and destination
volumes, removes snapshots on the destination volume that are newer than
the common snapshot and mounts the destination volume as a DP volume with
the common snapshot as the exported snapshot.
* For data protection (DP) relationships, takes a snapshot of the source
volume to capture the current image and transfers snapshots that are newer
than the common snapshot from the source volume to the destination volume.
For extended data protection (XDP) relationships, transfers snapshots
newer than the common snapshot according to the relationship policy, i.e.,
snapshots will match rules associated with the policy as defined by the
`snapmirror policy` commands. For relationships associated with
`snapmirror policy` of type ``_async-mirror_`` and ``_mirror-vault_`` the
`snapmirror resync` first takes a snapshot of the source volume and
includes it in the snapshots selected for transfer.
* For a SnapLock Compliance volume in an XDP relationship with SnapMirror
policy of type async-mirror, if SnapMirror resync operation detects data
divergence between the common snapshot and the AFS on the destination
volume, the resync operation preserves the data changes in a locked
snapshot for the duration of the current volume expiry time. If the volume
expiry time is in the past or has not been set, then the snapshot is
locked for a duration of 30 days. The common snapshot is also locked for
the same duration.



For Vserver SnapMirror relationships, a resync operation is typically
executed when the relationship is broken-off, the ``_subtype_`` of the
destination Vserver is ``_default_`` and the destination volumes are of
type read-write. Once the command is queued, the subtype of the
destination Vserver changes from ``_default_`` to ``_dp-destination_`` . A
successful resync operation also makes the destination Vserver's volumes
data protection volumes.

If the resync command is executed on a Vserver SnapMirror relationship,
and the corresponding source and destination Vservers have volumes with
volume level SnapMirror relationships, then the volume level SnapMirror
relationships will be converted to volumes under the Vserver SnapMirror
relationship. This conversion is supported only for source and destination
Vservers which have been transitioned from a 7-Mode vFiler into a C-Mode
Vserver. Some basic pre-requisites for the conversion are that the
destination Vserver should be in a ``_stopped_`` state and all the
destination Vserver volumes except the root volume should be in a volume
level SnapMirror relationship with volumes of the source Vserver. The
state of these volume level SnapMirror relationships should be
``_Snapmirrored_`` and status should be ``_Idle_`` .
```

```
When the destination Application Consistency Group relationship is broken-
off, the `snapmirror resync` changes the destination Application
Consistency Group to read-only and the relationship will be SnapMirrored.

`snapmirror resync` for a relationship with a policy of type ``_strict-
sync-mirror_`` or ``_sync-mirror_`` is typically executed in the following
case:

* The destination mirror is broken (that is, the destination volume is
read-write and no longer read-only). After the `snapmirror resync` command
completes, the destination volume changes to read-only and the
relationship to InSync.



NOTE: The `snapmirror resync` command is typically not required to return
a relationship that has fallen out of sync due to an error condition to
InSync because SnapMirror has Auto Resync for synchronous relationships.
When SnapMirror detects that the relationship has fallen out of sync for
any reason other than a xref:{relative_path}snapmirror-
quiesce.html[snapmirror quiesce] , xref:{relative_path}snapmirror-
break.html[snapmirror break] or xref:{relative_path}snapmirror-
delete.html[snapmirror delete] command was executed on the relationship,
it will automatically initiate a resync operation.



The default behavior of the `snapmirror resync` command for relationships
with a policy of type ``_strict-sync-mirror_`` or ``_sync-mirror_`` is
defined as follows:

* Creates a snapshot on the destination of the current image of the
destination file system. This snapshot becomes the exported snapshot for
the volume during the resync operation.
* Finds the most recent common snapshot between the source and destination
volumes. Performs a local rollback transfer to give the active file system
the same data as the common snapshot. It then loops through a sequence,
creating a snapshot on the source volume, transferring the data captured
in that snapshot, creating a snapshot of the data on the destination, and
repeating until the relationship is close to InSync. After the last
transfer, it enters cutover to bring the relationship to InSync.
* User-created snapshots are not replicated by a resync operation.
* At the conclusion of the resync operation, the exported snapshot on the
destination is removed and the client will then see the active file system
on the destination volume. The relationship will be InSync and periodic
creation of common snapshots will resume.
```

The `snapmirror resync` command supports an optional parameter ``_"preserve"_`` . The parameter ``_"preserve"_`` is only supported for extended data protection (XDP) relationships. It is not supported for relationships with a non-ONTAP endpoint. It is not supported for relationships with a policy of type ``_strict-sync-mirror_`` and ``_sync-mirror_`` . When used, the parameter ``_"preserve"_`` changes the behavior of the `snapmirror resync` command. The changed behavior of the command can be described as follows:

* Finds the most recent common snapshot between the source and destination volumes, preserves all snapshots on the destination volume that are newer than the common snapshot, and mounts the destination volume as a DP volume with the common snapshot as the exported snapshot.
* Performs a local rollback transfer to make a copy of the common snapshot on the destination volume and establish it as the latest snapshot on the destination volume. The command then transfers all snapshots that are newer than the common snapshot, from the source volume to the destination volume. The command only transfers snapshots that match the relationship's policy, i.e., snapshots will match rules associated with the policy as defined by the `snapmirror policy` commands.

If a SnapMirror relationship does not already exist, that is, the relationship was not created using the xref:{relative_path}snapmirror-create.html[snapmirror create] command, the `snapmirror resync` command will implicitly create the SnapMirror relationship, with the same behaviors as described for the xref:{relative_path}snapmirror-create.html[snapmirror create] command before resyncing it.

For Vservers, you must create SnapMirror relationships between Vservers by using the xref:{relative_path}snapmirror-create.html[snapmirror create] command before you run the `snapmirror resync` command. The `snapmirror resync` command does not implicitly create the relationship.

This command is supported for SnapMirror relationships with the field ``_"Relationship Capability"_`` showing as either ``_"8.2 and above"_`` or ``_"Pre 8.2"_`` in the output of the xref:{relative_path}snapmirror-show.html[snapmirror show] command.

For relationships with ``_"Relationship Capability"_`` of ``_"8.2 and above"_`` , you can track the progress of the operation using the xref:{relative_path}snapmirror-show.html[snapmirror show] command.

```
For relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_``
, a job will be spawned to operate on the SnapMirror relationship, and the
job id will be shown in the command output. The progress of the job can be
tracked using the xref:{relative_path}job-show.html[job show] and
xref:{relative_path}job-history-show.html[job history show] commands.

The `snapmirror resync` command fails if the destination volume does not
have a snapshot in common with the source volume.

The `snapmirror resync` command does not work on load-sharing mirrors.

The `snapmirror resync` command must be used from the destination Vserver
or cluster.

== Parameters

{ ``[-S, -source-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Source Path::
This parameter specifies the source endpoint of the SnapMirror
relationship in one of four path formats. The normal format includes the
names of the Vserver (vserver) and/or volume (volume). To support
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , a
format which also includes the name of the cluster (cluster) is provided.
The ``_"Pre 8.2"_`` format cannot be used when operating in a Vserver
context on relationships with ``_"Relationship Capability"_`` of ``_"8.2
and above"_`` . For SnapMirror relationships with an AltaVault source, the
source endpoint is specified in the form ``_hostip:/share/share-name_`` .
For SnapMirror relationships with a SolidFire source, the source endpoint
is specified in the form ``_hostip:/lun/name_`` .

| ``[-source-cluster <Cluster name>]`` - Source Cluster::
Specifies the source cluster of the SnapMirror relationship. If this
parameter is specified, the `-source-vserver` and `-source-volume`
parameters must also be specified. This parameter is only applicable for
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .
This parameter cannot be specified when operating in a Vserver context on
relationships with ``_"Relationship Capability"_`` of ``_"8.2 and
above"_`` .

``[-source-vserver <vserver name>]`` - Source Vserver::
Specifies the source Vserver of the SnapMirror relationship. For
relationships with volumes as endpoints, if this parameter is specified,
parameters `-source-volume` and for relationships with ``_"Relationship
Capability"_`` of ``_"Pre 8.2"_`` , `-source-cluster` must also be
specified. This parameter is not supported for relationships with non-
```

```
ONTAP source endpoints.

``[-source-volume <volume name>]`` - Source Volume }::
Specifies the source volume of the SnapMirror relationship. If this
parameter is specified, parameters `-source-vserver` and for relationships
with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-source-
cluster` must also be specified. This parameter is not supported for
relationships with non-ONTAP source endpoints.

{ ``-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}`` - Destination
Path::
This parameter specifies the destination endpoint of the SnapMirror
relationship in one of four path formats. The normal format includes the
names of the Vserver (vserver) and/or volume (volume). To support
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , a
format which also includes the name of the cluster (cluster) is provided.
The ``_"Pre 8.2"_`` format cannot be used when operating in a Vserver
context on relationships with ``_"Relationship Capability"_`` of ``_"8.2
and above"_`` . For SnapMirror relationships with an AltaVault
destinations, the destination endpoint is specified in the form
``_hostip:/share/share-name_`` . For SnapMirror relationships with a
SolidFire destinations, the destination endpoint is specified in the form
``_hostip:/lun/name_`` .

| ``[-destination-cluster <Cluster name>]`` - Destination Cluster::
Specifies the destination cluster of the SnapMirror relationship. If this
parameter is specified, parameters `-destination-vserver` and `-
destination-volume` must also be specified. This parameter is only
applicable for relationships with ``_"Relationship Capability"_`` of
``_"Pre 8.2"_`` . This parameter cannot be specified when operating in a
Vserver context on relationships with ``_"Relationship Capability"_`` of
``_"8.2 and above"_`` .

``-destination-vserver <vserver name>`` - Destination Vserver::
Specifies the destination Vserver of the SnapMirror relationship. For
relationships with volumes as endpoints, if this parameter is specified,
parameters `-destination-volume` and for relationships with
``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-destination-
cluster` must also be specified. This parameter is not supported for
relationships with non-ONTAP destination endpoints.

``-destination-volume <volume name>`` - Destination Volume }::
Specifies the destination volume of the SnapMirror relationship. If this
parameter is specified, parameters `-destination-vserver` and for
```

```
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-
destination-cluster` must also be specified. This parameter is not
supported for relationships with non-ONTAP destination endpoints.

``[-s, -source-snapshot <text>]`` - Source Snapshot::
This optional parameter specifies a snapshot to transfer. The default
behavior for a data protection relationship with a read-write source is
that ONTAP creates a new snapshot and uses it as the basis for determining
what data is replicated; with this option, the specified snapshot will be
used instead. The default behavior for an extended data protection
relationship depends on the relationship's policy type. For a data
protection relationship, the specified snapshot must be newer than the
latest common snapshot. For an extended data protection relationship, the
specified snapshot can be newer or older than the common snapshot. This
parameter is not supported for relationships with ``_"Relationship
Capability"_`` of ``_"Pre 8.2"_`` .

``[-type <snapmirrorType>]`` - Snapmirror Relationship Type::
Specifies the type of SnapMirror relationship if a relationship is
implicitly created. This parameter is the same as the one used in the
xref:{relative_path}snapmirror-create.html[snapmirror create] command.

``[-policy <sm_policy>]`` - SnapMirror Policy::
This optional parameter designates the name of the SnapMirror policy which
is associated with the SnapMirror relationship. If you do not designate a
policy, the current policy will be retained. This parameter is not
applicable to relationships with ``_"Relationship Capability"_`` of
``_"Pre 8.2"_`` .
+
NOTE: You define and name a policy using the
xref:{relative_path}snapmirror-policy-create.html[snapmirror policy
create] command.
+



``[-f, -force <true>]`` - Force::
If this parameter is specified, the command proceeds without prompting for
confirmation.

``[-k, -throttle <throttleType>]`` - Throttle (KB/sec)::
This optional parameter limits the network bandwidth used for the resync
transfer. It sets the maximum rate (in Kbytes/sec) at which data can be
transferred during the operation. If this parameter is not specified, the
throttle value configured for the relationship with the
xref:{relative_path}snapmirror-create.html[snapmirror create] or
xref:{relative_path}snapmirror-modify.html[snapmirror modify] command will
```

be used. To fully use the network bandwidth available, set the throttle value to ``_unlimited_`` or ``_0_`` . The minimum throttle value is four Kbytes/sec, so if you specify a throttle value between ``_1_`` and ``_4_`` , it will be treated as if you specified ``_4_`` . For FlexGroup and Application Consistency Group relationships, the throttle value is applied individually to each constituent relationship. For relationships with a policy of type ``_strict-sync-mirror_`` or ``_sync-mirror_`` , the throttle value is applicable only for the asynchronous transfers of the resync operation. The `-throttle` parameter does not affect load-sharing transfers and transfers for other relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` confined to a single cluster.

``[-transfer-priority {low|normal}]`` - Transfer Priority::
This optional parameter specifies the priority at which the transfer runs. The default value for this parameter is the value in the SnapMirror policy associated with this relationship. This parameter is not applicable to relationships with a ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .

``[-preserve <true>]`` - Preserve::
This parameter is only supported for extended data protection (XDP) relationships with policies of type ``_vault_`` , and ``_mirror-vault_`` . It is not supported for relationships with a policy of type ``_async-mirror_`` and data protection and load-sharing relationships. This parameter is not supported for relationships with non-ONTAP endpoints. It is not supported for relationships with a policy of type ``_strict-sync-mirror_`` and ``_sync-mirror_`` . When specified, it changes the behavior of the `snapmirror resync` command to preserve snapshots on the destination volume that are newer than the latest common snapshot. This parameter is not supported for relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .

``[-quick-resync <true>]`` - Quick Resync::
This parameter is only supported for extended data protection (XDP) relationships. This parameter is not supported for relationships with non-ONTAP endpoints. It is not supported for relationships with a policy of type ``_strict-sync-mirror_`` and ``_sync-mirror_`` . Specifying this optional parameter reduces the resync time because the resync does not incur storage efficiency overhead before the transfer of new data. Specifying this parameter is recommended if the source of the resync does not have volume efficiency enabled or if reducing resync time is more important than preserving all possible storage efficiency. When this parameter is specified, resync does not preserve the storage efficiency of the new data with existing data over the wire and on the destination.

``[-cg-item-mappings {<source volume>:@<destination volume>|<source item path>:@<destination item path>}]`` - Consistency Group Item Mappings::

```
This optional parameter specifies a list of the consistency group (CG)
item mappings. It is considered only if the supplied source and
destination path values denote valid CG paths. For CG level relationships,
this parameter must be specified. The value for this parameter must
contain mappings of one or more pairs of constituent source and
destination volumes of the form ``_srcvol1:@dstvol1,..._`` .

``[-is-auto-expand-enabled {true|false}]`` - Is Auto Expand Enabled::
This optional parameter specifies whether or not a FlexGroup SnapMirror
relationship and its destination FlexGroup should be auto-expanded if the
source FlexGroup is expanded. This parameter is supported only for
FlexGroup SnapMirror relationships. If this resync is creating a new
SnapMirror relationship, the default value is true. If it is not creating
a new relationship, if a value is specified, it must match the current
value for the existing relationship. If the parameter is not specified,
the existing value will be retained.

``[-backoff-level {high|medium|none}]`` - SM Backoff Level due to Client
Ops::
This optional parameter specifies the SnapMirror backoff level due to
client ops. This parameter is supported only for FlexVol SnapMirror
relationships. The default value is `high` .

``[-w, -foreground <true>]`` - Foreground Process::
This specifies whether the operation runs as a foreground process. If this
parameter is specified, the default setting is ``_true_`` (the operation
runs in the foreground). When set to ``_true_`` , the command will not
return until the process completes. This parameter is only applicable to
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .

== Examples

To reestablish mirroring for the destination endpoint
vs2.example.com:dept_mkt_mirror that has been previously broken off with
the xref:{relative_path}snapmirror-break.html[snapmirror break] command,
type the following command:
```

vs2.example.com::> snapmirror resync -destination-path
vs2.example.com:dept_mkt_dp_mirror

```
For relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_``
, to reestablish mirroring for the destination endpoint
``_cluster2://vs2.example.com/dept_mkt_mirror_`` that has been previously
broken off with the xref:{relative_path}snapmirror-break.html[snapmirror
break] command, type the following command:
```

cluster2::> snapmirror resync -destination-path
cluster2//vs2.example.com/dept_mkt_dp_mirror

```
To create a SnapMirror relationship and reestablish mirroring between the
destination endpoint named ``_vs2.example.com:dept_eng_dp_mirror2_`` and
the source endpoint named ``_vs1.example.com:dept_eng_`` , type the
following command:
```

vs2.example.com::> snapmirror resync -destination-path
vs2.example.com:dept_eng_dp_mirror2
-source-path vs1.example.com:dept_eng

```
To create a SnapMirror relationship and reestablish mirroring between the
destination endpoint named
``_cluster2://vs2.example.com/dept_eng_dp_mirror2_`` and the source
endpoint named ``_cluster1://vs1.example.com/dept_eng_`` when the source
cluster is running ONTAP 8.1 software, type the following command:
```

cluster2::> snapmirror resync -destination-path
cluster2://vs2.example.com/dept_eng_dp_mirror2
-source-path cluster1://vs1.example.com/dept_eng

```
To create and reestablish an extended data protection (XDP) relationship
between the ONTAP source endpoint ``_vs1.example.com:data_ontap_vol_`` ,
and the non-ONTAP (for example, AltaVault) destination endpoint
``_10.0.0.11:/share/share1_`` , and start the initial transfer, type the
following command:
```

vs1.example.com::> snapmirror resync -destination-path
10.0.0.11:/share/share1
-source-path vs1.example.com:data_ontap_vol -type XDP

```
To reestablish mirroring for the destination endpoint
``_dvs1.example.com:_`` of a Vserver relationship that has been previously
broken off with the xref:{relative_path}snapmirror-break.html[snapmirror
break] command, type the following command:
```

cluster2::> snapmirror resync -destination-path
dvs1.example.com:

```
To resynchronize the SnapMirror active sync relationship with the source
Consistency Group ``_cg_src_`` in Vserver ``_vs1.example.com_`` and the
destination Consistency Group ``_cg_dst_`` in Vserver
``_vs2.example.com_`` , type the following command from the destination
cluster:
```

destination::> snapmirror resync -destination-path
vs2.example.com:/cg/cg_dst

```
To resynchronize an SnapMirror active sync relationship with the following
attributes:

* It is between the source Consistency Group ``_cg_src_`` in Vserver
``_vs1.example.com_`` , and the destination Consistency Group ``_cg_dst_``
in Vserver ``_vs2.example.com_`` .
* It has item mappings between volumes ``_srcvol1_`` and ``_srcvol2_`` and
volumes ``_dstvol1_`` and ``_dstvol2_`` .

Type the following command from the destination cluster:
```

destination::> snapmirror resync -destination-path
vs2.example.com:/cg/cg_dst -source-path
vs1.example.com:/cg/cg_src
-cg-item-mappings srcvol1:@dstvol1,srcvol2:@dstvol2

```
To reestablish mirroring to the destination Application Consistency Group
``_app_cg_dst_`` in Vserver ``_vs2.example.com_`` that has been previously
broken off with the xref:{relative_path}snapmirror-break.html[snapmirror
break] command, type the following command:
```

vs2.example.com::> snapmirror resync -destination-path
vs2.example.com:/cg/app_cg_dst

```
== Related Links

* xref:{relative_path}snapmirror-update.html[snapmirror update]
* xref:{relative_path}snapmirror-quiesce.html[snapmirror quiesce]
* xref:{relative_path}snapmirror-break.html[snapmirror break]
* xref:{relative_path}snapmirror-delete.html[snapmirror delete]
* xref:{relative_path}snapmirror-create.html[snapmirror create]
* xref:{relative_path}snapmirror-show.html[snapmirror show]
* xref:{relative_path}job-show.html[job show]
* xref:{relative_path}job-history-show.html[job history show]
```

```
* xref:{relative_path}snapmirror-policy-create.html[snapmirror policy
create]
* xref:{relative_path}snapmirror-modify.html[snapmirror modify]




[[IDb807e39db7c3c3c95bdd4730a18d50e9]]
= snapmirror set-options
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display/Set SnapMirror options

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror set-options` command can be used to display or set
SnapMirror options.

== Parameters

``[-dp-source-xfer-reserve-pct {0|25%|50%|75%|100%}]`` - Percentage
Reserved for DP Source Transfers::
Specifies the percentage of maximum allowed concurrent transfers reserved
for source DP transfers

``[-xdp-source-xfer-reserve-pct {0|25%|50%|75%|100%}]`` - Percentage
Reserved for XDP Source Transfers::
Specifies the percentage of maximum allowed concurrent transfers reserved
for source XDP transfers

``[-dp-destination-xfer-reserve-pct {0|25%|50%|75%|100%}]`` - Percentage
Reserved for DP Destination Transfers::
Specifies the percentage of maximum allowed concurrent transfers reserved
for destination DP transfers

``[-xdp-destination-xfer-reserve-pct {0|25%|50%|75%|100%}]`` - Percentage
Reserved for XDP Destination Transfers::
Specifies the percentage of maximum allowed concurrent transfers reserved
for destination XDP transfers
```

```
== Examples

The following example displays SnapMirror options:
```

cluster1::> snapmirror set-options
Percentage Reserved for DP Source Transfers: 0
Percentage Reserved for XDP Source Transfers: 0
Percentage Reserved for DP Destination Transfers: 0
Percentage Reserved for XDP Destination Transfers: 0

cluster1::> snapmirror set-options -dp-source-xfer-reserve-pct 25
-xdp-source-xfer-reserve-pct 50 -dp-destination-xfer
reserve-pct 0 -xdp-destination-xfer-reserve-pct 50

cluster1::> snapmirror set-options
Percentage Reserved for DP Source Transfers: 25
Percentage Reserved for XDP Source Transfers: 50
Percentage Reserved for DP Destination Transfers: 0
Percentage Reserved for XDP Destination Transfers: 50

```
[[ID88e530d88e2b37a33ce67709d723f9db]]
= snapmirror set-preferred-cluster
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set this cluster as the preferred cluster

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror set-preferred-cluster` sets the local cluster as the
preferred cluster for SnapMirror active sync relationships of policy-type
'automated-failover-duplex'. Note that the 'preferred-cluster' will be the
one serving IO upon scenarios like a network partition between primary and
secondary.

You should execute the `snapmirror set-preferred-cluster` command from the
destination cluster.

Make sure the relationship status is InSync and the ONTAP Mediator is
configured, connected, and in quorum before using this command.

== Parameters

``-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}`` - Destination
Path::
Specifies the destination endpoint of the SnapMirror relationship.

== Examples

To perform a set-preferred-cluster operation of an SnapMirror active sync
relationship with source Consistency Group ``_cg_src_`` in Vserver
``_vs1.example.com_`` , and the destination Consistency Group ``_cg_dst_``
in Vserver ``_vs2.example.com_`` , type the following command from the
destination cluster:
```

destination::> snapmirror set-preferred-cluster -destination-path
vs2.example.com:/cg/cg_dst

```
Note: The SnapMirror set-preferred-cluster operation internally triggers a
planned failover operation. The SnapMirror active sync relationship
information in the xref:{relative_path}snapmirror-show.html[snapmirror
show] command will now be available on the new destination that is Vserver
``_vs2.example.com_`` with ``_cg_dst_`` as the new source and ``_cg_src_``
as the new destination Consistency Groups after the SnapMirror set-
preferred-cluster operation is completed.

== Related Links

* xref:{relative_path}snapmirror-show.html[snapmirror show]




[[IDc50f1e64a1324ddfd7697b595ca3d514]]
= snapmirror show-history
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays history of SnapMirror operations.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror show-history` command displays the history of the last 24
hours of SnapMirror operations. This command is not supported for
relationships with non-ONTAP endpoints.

By default, the command displays the following information:

* Destination Path
* Source Path
* Operation
* Start Time
* End Time
* Result



The `snapmirror show-history` command displays - in reverse chronological
order - the history of completed SnapMirror operations whose destination
```

```
endpoints are in the current Vserver for Vserver administrators, or the
current cluster for cluster administrators. This command does not return
information on the following operations and relationships:

* Operations that happened prior to installing ONTAP 8.3.
* Relationships with the ``_"Relationship Capability"_`` field, as shown
in the output of the SnapMirror show command, set to ``_"Pre 8.2"_`` .
* Operations on FlexGroup relationships that happened prior to installing
ONTAP 9.5.
* Operations on FlexGroup constituent relationships.



The `-instance` parameter displays the following detailed information:
```

```
       Destination Path: Path of the destination endpoint.
            Source Path: Path of the source endpoint.
        Relationship ID: The unique identifier of the relationship.
                         This parameter is not supported for
                         Vserver SnapMirror relationships.
Relationship Group Type: For FlexVol relationships, specifies the
                         type of the group relationship that
                         includes this FlexVol. For group
                         relationships, specifies the type of the
                         group relationship. Can be one of the
                         following:
                         - none: No group relationship.
                         - vserver: Vserver relationship.
                         - flexgroup: FlexGroup relationship.
              Operation: Type of the operation.
                         Can be one of the following:
                         - create
                         - modify
                         - quiesce
                         - resume
                         - delete
                         - initialize
                         - manual update
                         - scheduled update
                         - break
                         - resync
                         - abort
                         - restore
           Operation ID: The unique identifier of the operation.
             Start Time: Timestamp of the start of the operation.
               End Time: Timestamp of the end of the operation.
                 Result: Result of the SnapMirror operation.
                         Can be one of the following:
                         - success
                         - failure
          Transfer Size: Total amount of data transferred during the
                         SnapMirror operation.
 Additional Information: A message describing the cause of the
                         failure or additional information about a
                         successful operation, such as if a checkpoint
                         was cleared as part of an abort operation.
```

```
== 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.

{ ``[-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Destination
Path::
Select SnapMirror operations that have a matching destination path name.

| ``[-destination-vserver <vserver name>]`` - Destination Vserver::
Select SnapMirror operations that have a matching destination Vserver
name.

``[-destination-volume <volume name>]`` - Destination Volume }::
Select SnapMirror operations that have a matching destination volume name.

``[-operation-id <UUID>]`` - Operation ID::
Select SnapMirror operations that have a matching operation ID.

``[-source-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Source Path::
Select SnapMirror operations that have a matching source path name.

``[-source-vserver <vserver name>]`` - Source Vserver::
Select SnapMirror operations that have a matching source Vserver name.

``[-source-volume <volume name>]`` - Source Volume::
Select SnapMirror operations that have a matching source volume name.

``[-operation-type {create|modify|quiesce|resume|delete|initialize|manual-
update|scheduled-
update|break|resync|abort|restore|failover|rsetup|baseline-gather}]`` -
Operation Type::
Select SnapMirror operations that have a matching operation type. Possible
values are:
+
  * create
  * modify
  * quiesce
```

```
  * resume
  * delete
  * initialize
  * manual-update
  * scheduled-update
  * break
  * resync
  * abort
  * restore



``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Start Time::
Select SnapMirror operations that have a matching start time.

``[-end-time <MM/DD/YYYY HH:MM:SS>]`` - End Time::
Select SnapMirror operations that have a matching end time.

``[-relationship-id <UUID>]`` - Relationship ID::
Select SnapMirror operations that have a matching relationship ID.

``[-relationship-group-type
{none|vserver|consistencygroup|flexgroup|vserverflexgroup|vserverconsisten
cygroup}]`` - Relationship Group Type::
Select SnapMirror relationships that have a matching relationship group
type. Possible values are:
+
  * ``_none_``
  * ``_vserver_``
  * ``_flexgroup_``



``[-result {success|failure}]`` - Result of the Operation::
Select SnapMirror operations that have a matching result. Possible values
are:
+
  * success
  * failure



``[-transfer-size {<integer>[KB|MB|GB|TB|PB]}]`` - Transfer Size::
Select SnapMirror operations that have a matching transfer size.

``[-additional-info <text>]`` - Additional Information::
Select SnapMirror operations that have matching additional information.
```

```
``[-max-rows-per-relationship <integer>]`` - Maximum Number of Rows per
Relationship::
Select a matching number of SnapMirror operations per relationship.

``[-expand <true>]`` - Show Constituents of the Group.::
Select SnapMirror operations on relationships that are constituents and
non-constituents of a group.

== Examples

The example below displays summary information for all SnapMirror
operations on relationships with destination endpoints in the current
cluster:
```

cluster2::> snapmirror show-history
Destination Source Start End
Path Path Operation Time Time Result
---------- ---------- --------- ---------- ---------- -------
dvs2: vs2: create 8/15/2013 08:22:23
8/15/2013 08:22:25
success
vs1:vol1 vs1:aggr1 manual-update
8/15/2013 08:22:44
8/15/2013 08:22:44
failure
vs1:vol1 vs1:aggr1 initialize
8/15/2013 08:22:25
8/15/2013 08:22:26
success
vs1:vol1 vs1:aggr1 create 8/15/2013 08:22:15
8/15/2013 08:22:16
success
vs1:vol2 vs1:aggr1 initialize
8/15/2013 08:23:23
8/15/2013 08:23:23
failure
vs1:vol2 vs1:aggr1 create 8/15/2013 08:23:10
8/15/2013 08:23:10
success
6 entries were displayed.

```
The example below displays detailed information for the SnapMirror
operation with operation ID ``_dc158715-0583-11e3-89bd-123478563412_``
```

cluster2::> snapmirror show-history -operation-id dc158715-0583-11e3-89bd-123478563412
Destination Path: vs1:vol1
Source Path: vs1:aggr1
Relationship ID: cb3d30a0-0583-11e3-89bd-123478563412

Relationship Group Type: none
Operation: manual-update
Operation ID: dc158715-0583-11e3-89bd-123478563412
Start Time: 8/15/2013 08:22:44
End Time: 8/15/2013 08:22:44
Result: failure
Transfer Size: -
Additional Information: Volume vs1:vol1 is restricted. Use the command "volume online" to bring the volume online.

```
The example below displays detailed information for all SnapMirror
operations on relationships with the ``_Result_`` of ``_"success"_`` and
whose destination endpoints are in the current cluster.
```

cluster2::> snapmirror show-history -result success -instance
Destination Path: vs1:vol1
Source Path: vs1:aggr1
Relationship ID: cb3d30a0-0583-11e3-89bd-123478563412
Relationship Group Type: none
Operation: initialize
Operation ID: d03ce1db-0583-11e3-89bd-123478563412
Start Time: 8/15/2013 08:22:25
End Time: 8/15/2013 08:22:26
Result: success
Transfer Size: 1.09MB
Additional Information: -
Destination Path: vs1:vol1
Source Path: vs1:aggr1
Relationship ID: cb3d30a0-0583-11e3-89bd-123478563412
Relationship Group Type: none
Operation: create
Operation ID: cb3d305d-0583-11e3-89bd-123478563412
Start Time: 8/15/2013 08:22:15
End Time: 8/15/2013 08:22:16
Result: success
Transfer Size: -
Additional Information: -
Destination Path: vs1:vol2
Source Path: vs1:aggr1
Relationship ID: eb92c549-0583-11e3-89bd-123478563412
Relationship Group Type: none
Operation: create
Operation ID: eb92c506-0583-11e3-89bd-123478563412
Start Time: 8/15/2013 08:23:10
End Time: 8/15/2013 08:23:10
Result: success
Transfer Size: -
Additional Information: -

3 entries were displayed.

```
The example below displays summary information for all SnapMirror
operations    on relationships with ``_max-rows-per-relationship_`` of
``_1_`` and whose destination endpoints are in the current cluster.
```

cluster2::> snapmirror show-history -max-rows-per-relationship 1

Destination Source Start End
Path Path Operation Time Time Result
---------- ---------- --------- ---------- ---------- -------
vs1:vol1 vs1:aggr1 manual-update
8/15/2013 08:22:44
8/15/2013 08:22:44
failure
vs1:vol2 vs1:aggr1 initialize
8/15/2013 08:23:23
8/15/2013 08:23:23
failure
2 entries were displayed.

```
The example below displays summary information for all SnapMirror
operations on relationships between Consistency Groups with destination
endpoints in the current cluster:
```

cluster2::> snapmirror show-history

Destination Source Start End
Path Path Operation Time Time Result
---------- ---------- --------- ---------- ---------- -------
svm2:/cg/cg1_dst
svm1:/cg/cg1
manual-update
10/10/2024 01:01:17
10/10/2024 01:10:40
success
svm2:/cg/cg1_dst
svm1:/cg/cg1
initialize
10/9/2024 13:53:16
10/9/2024 13:55:49
success
svm2:/cg/cg1_dst
svm1:/cg/cg1
create 10/9/2024 13:52:40
10/9/2024 13:52:44
success
3 entries were displayed.

```
[[ID9175a09912ecc96285820098ec4ece02]]
= snapmirror show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of SnapMirror relationships

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror show` command displays information associated with
SnapMirror relationships. By default, the command displays the following
information:

* Source path
* Relationship Type
* Destination Path
* Mirror State
* Relationship Status
* Total Progress
* Healthy
* Progress Last Updated



For backward compatibility with ONTAP 8.1, and to accommodate load-sharing
relationships which are only supported in a ONTAP 8.1 compatible way,
SnapMirror relationships, which match one of the following conditions are
managed as on ONTAP 8.1: (1) The relationship is of type load-sharing; (2)
The source endpoint of the relationship is on a remote ONTAP 8.1 cluster;
(3) The local cluster was upgraded from ONTAP 8.1, the relationship was
created before the upgrade, and the relationship has not yet been
converted to one with ONTAP 8.2 capabilities. These relationships have the
same limitations as on ONTAP 8.1. Especially, they support the same set of
information fields. The ``_"Relationship Capability"_`` field is set to
``_"Pre 8.2"_`` for these relationships.

The `snapmirror show` command displays information for SnapMirror
relationships whose destination endpoints are in the current Vserver if
you are in a Vserver context, or in the current cluster if you are in a
cluster context, or on a non-ONTAP endpoint that supports SnapMirror (for
example, AltaVault). For backward compatibility with ONTAP 8.1, the
command also displays information for SnapMirror relationships with the
```

```
``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , and whose source
endpoints are in the current Vserver or cluster, and destination endpoints
are in different Vservers or clusters. You must use the
xref:{relative_path}snapmirror-list-destinations.html[snapmirror list-
destinations] command to display information for SnapMirror relationships
whose source endpoints are in the current Vserver or current cluster.

Some of the SnapMirror relationship information is cached. The `snapmirror
show` command only returns the cached information, therefore there is a
delay after the information is changed before it is reflected in the
`snapmirror show` output. Other information, such as progress metrics
during a transfer, is only updated periodically and can be very delayed in
the `snapmirror show` output.

The `-instance` and `-fields` parameters are mutually exclusive and select
the information fields that are displayed. The other parameters to the
`snapmirror show` command select the SnapMirror relationships for which
information is displayed. The `-instance` displays detailed information
fields including:
```

Source Path: Path of the source endpoint.
Destination Path: Path of the destination endpoint.
Relationship Type: Type of the SnapMirror relationship. Can be
one of the following:
- DP: Data protection relationship.
- LS: Load-sharing relationship.
- XDP: Extended data protection relationship.
- RST: Temporary relationship created
during a restore operation, and
deleted if the operation completes
successfully.
- TDP: 7-mode to ONTAP
transition data protection
relationship.
Relationship Group Type: For FlexVol relationships, specifies the
type of the group relationship that includes
this FlexVol. For group relationships,
specifies the type of the group
relationship. Can be one of the following:
- none: No group relationship.
- vserver: Vserver relationship.
- flexgroup: FlexGroup relationship.
- consistencygroup: Consistency Group
relationship.

```

Only for relationships with
"Relationship Capability"
```

of "8.2 and above" .
Relationship Status: Status of the SnapMirror relationship.
Can be one of the following:
- Idle: No transfer operation is in
progress and future transfers are
not disabled.
- Queued: A transfer operation has been
accepted and queued in the system,
and future transfers are not
disabled.
- Transferring: A transfer operation is in
progress and future transfers are
not disabled.
- Preparing: Pre-transfer phase for
Vault incremental transfers.
For Vault relationships only.
- Finalizing: Post-transfer phase for
Vault incremental transfers.
Network traffic will be low as
processing is primarily on the
destination volume.
For Vault relationships only.
- Aborting: A transfer abort operation
that might include the removal of
the checkpoint is underway. Future
transfers are not disabled. Only
for relationships with
"Relationship Capability"
of "8.2 and above" .
- Quiesced: No transfer operation is in
progress and future transfers are
disabled.
- Quiescing: A transfer operation is in
progress and future transfers
are disabled.
- Checking: Destination volume is
undergoing a diagnostic check,
no transfer is in progress, and
future transfers are not disabled.
Only for relationships with
"Relationship Capability"
of "Pre 8.2" .
- Breaking: The SnapMirror relationship is
being broken off and no transfer is
in progress.

The following values are only applicable to
relationships with a policy of type
sync-mirror, strict-sync-mirror,
automated-failover, or
automated-failover-duplex:

- OutOfSync: The SnapMirror relationship is not InSync and no async transfer operation is in progress.
- Transitioning: The SnapMirror relationship is transitioning to InSync.
- InSync: The SnapMirror relationship is InSync.

The following values are only applicable to relationships with a policy of type automated-failover:
- Modifying: One or more properties of the SnapMirror relationship are being modified.

Mirror State: State of the destination volume. Can be one of the following:
- Uninitialized: Destination volume has not been initialized.
- Snapmirrored: Destination volume has been initialized and is ready to receive SnapMirror updates.
- Broken-off: Destination volume is RW and snapshots are present.

```
 ----
The following values are applicable only to
                              the relationships with a policy of type
            automated-failover or
            automated-failover-duplex:
                              - Expanding: Adding new volumes to
                                           the SnapMirror relationship is
                                           in progress.
 ----
Healthy: Condition of the relationship. Can be one of
                              the following:
                              - true: The SnapMirror relationship is
                                      healthy. It has not missed a
                                      scheduled transfer, or experienced
                                      a manual update failure.
                              - false: The SnapMirror relationship is not
                                     healthy. It has missed a scheduled
                                     transfer, or has experienced a manual
                                     update failure.
```

```
            Unhealthy Reason: Reason the SnapMirror relationship is not
                              healthy. Only for relationships with
                               "Relationship Capability"
                              of  "8.2 and above"
             Newest Snapshot: Name of the newest snapshot on the
                              destination volume.
   Newest Snapshot Timestamp: Timestamp of the newest snapshot.
           Exported Snapshot: Name of the exported snapshot on the
            destination volume. This field is not applicable
              to the policy of type
             automated-failover-duplex .
 Exported Snapshot Timestamp: Timestamp of the exported snapshot.
                   Lag Time: Time since the exported snapshot
                              was created. It is displayed in the
                              format: hours:minutes:seconds.
                              Only for relationships with
                               "Relationship Capability"
                              of  "8.2 and above" .
               Transfer Type: Type of the current transfer operation.
                              Can be one of the following:
                              - initialize
                              - update
                              - resync
                              - restore
                              Only for relationships with
                               "Relationship Capability"
                              of  "8.2 and above" .
           Transfer Snapshot: Name of the snapshot being transferred.
           Snapshot Progress: Amount of data transferred for the transfer
                              snapshot. This parameter is not supported
                              for SnapMirror FlexGroup relationships, but
                              it is supported for FlexGroup constituent
                              relationships.
              Total Progress: Total amount of data transferred for the
                              current transfer operation.
   Percent Complete for Current Status: Percent complete for the current
value
                              of status. This field is only valid when the
             "Relationship Status"  is
             "Finalizing" .
   Network Compression Ratio: The compression ratio achieved for the data
                              sent over the wire as a part of the current
                              transfer operation. The ratio is not
                              maintained across checkpoint restarts. If
                              network compression is disabled for the
                              transfer, the ratio will be set to 1:1.
```

```
                                Only for relationships with
                                 "Relationship Capability"
                                of  "8.2 and above" .
                                This parameter is not supported for Vserver
                                or FlexGroup SnapMirror relationships, but
                                it is supported for FlexGroup constituent
                                relationships.
         Snapshot Checkpoint: The amount of data transferred as recorded
in
                                the restart checkpoint of the current or
most
                                recent transfer snapshot. If a restart
                                checkpoint is present the next transfer will
                                continue from the checkpoint. This parameter
                                is not supported for SnapMirror FlexGroup
                                relationships, but it is supported for
                                FlexGroup constituent relationships.
               Transfer Error: Possible transient error condition if any,
                                encountered by the current transfer
                                operation.
                                Only for relationships with
                                 "Relationship Capability"
                                of  "8.2 and above".
            Current Throttle: The maximum transfer rate in Kilobytes
                                per second, used for the current transfer
                                between clusters.
                                Only for relationships with
                                 "Relationship Capability"
                                of  "8.2 and above" .
   Current Transfer Priority: Priority assigned to the current transfer.
                                Possible values are:
                                - low
                                - normal
                                Only for relationships with
                                 "Relationship Capability"
                                of  "8.2 and above" .
         Last Transfer Type: Type of the previous transfer operation:
                                - initialize
                                - update
                                - resync
                                - restore
                                Only for relationships with
                                 "Relationship Capability"
                                of  "8.2 and above" .
         Last Transfer Size: Total amount of data transferred during the
                                previous transfer operation if it was
```

```
                                 successful.
                                 Only for relationships with
                                  "Relationship Capability"
                                 of  "8.2 and above" .
                                 This parameter is not supported for
                                 SnapMirror FlexGroup relationships, but it
                                 is supported for FlexGroup constituent
                                 relationships.
Last Transfer Network Compression Ratio: The compression ratio achieved
                                 for the data sent over the wire as a part of
                                 the previous transfer operation. If network
                                 compression was disabled for the transfer,
                                 the ratio will be set to 1:1.
                                 Only for relationships with
                                  "Relationship Capability"
                                 of  "8.2 and above" .
                                 This parameter is not supported for Vserver
                                 or FlexGroup SnapMirror relationships, but
it
                                 is supported for FlexGroup constituent
                                 relationships.
      Last Transfer Duration: Duration of the previous transfer
                                 operation if it was successful.
                                 Only for relationships with
                                  "Relationship Capability"
                                 of  "8.2 and above" .
          Last Transfer From: Source endpoint of the previous transfer
                                 operation.
                                 Only for relationships with
                                  "Relationship Capability"
                                 of  "8.2 and above" .
 Last Transfer End Timestamp: Timestamp of the end of the previous
                                 transfer operation.
                                 Only for relationships with
                                  "Relationship Capability"
                                 of  "8.2 and above" .
         Last Transfer Error: Cause of the failure of the previous
                                 transfer operation.
                                 Only for relationships with
                                  "Relationship Capability"
                                 of  "8.2 and above" .
     Relationship Capability: Management and control compatibility:
                                 - "Pre 8.2": Management and control of
                                   the relationship is compatible with
                                   ONTAP 8.1.
                                 - "8.2 and above": Full support of
```

```
                                 ONTAP 8.2 or later SnapMirror
                                 relationship management and control.
                               This parameter is not supported for Vserver
                               SnapMirror relationships.
             Relationship ID: The unique identifier of the relationship.
                               Only for relationships with
                                "Relationship Capability"
                               of  "8.2 and above" .
                               This parameter is not supported for Vserver
                               SnapMirror relationships.
        Current Operation ID: Operation unique identifier of the currently
                               executing SnapMirror operation.
                               Only for relationships with
                                "Relationship Capability"
                               of  "8.2 and above" .
           Throttle (KB/sec): Configured maximum transfer rate for
                               cross-cluster transfers.
      SnapMirror Policy Type: Type of the SnapMirror policy associated
                               with the relationship. Can be one of the
                               following:
                               - async-mirror
                               - vault
                               - mirror-vault
```

```
```

Refer to the man page for the
snapmirror policy create command
for a description of what these types mean.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
SnapMirror Policy: Name of the SnapMirror policy associated
with the relationship.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
SnapMirror Schedule: Name of the schedule (empty if there is
no schedule) associated with the
relationship.
Tries Limit: Maximum number of times a transfer will be
tried.
Only for relationships with
"Relationship Capability"
of "Pre 8.2" .
This parameter is not supported for Vserver
SnapMirror relationships.
Destination Volume Node: Node which owns the destination volume
of the relationship. For FlexGroup

relationships it is the node which owns
the root constituent destination volume.
For object store relationships it is the
node which owns the source volume.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Identity Preserve Vserver DR: Whether or not the identity of the source
Vserver is replicated to the destination
Vserver. Can be:
- true: Source Vserver’s configuration will
additionally be replicated to the
destination, along with the
Vserver’s volumes and RBAC
configuration.
- false: Only volumes and RBAC configuration
of the source Vserver is replicated
to the destination.
This parameter is supported only for Vserver
SnapMirror relationships.
Volume MSIDs Preserved: Whether or not the MSIDs of the source
volumes are retained while creating
destination volumes.
Can be:
- true: MSIDs of source Vserver volumes and
destination Vserver volumes match.
- false: MSIDs of source Vserver volumes and
destination Vserver volumes do not
match.
This parameter is supported only for Vserver
SnapMirror relationships.
Is Auto Expand Enabled: Whether or not the auto expand is enabled.
Can be:
- true: Auto Expand is enabled.
- false: Auto Expand is disabled.
This parameter is supported only for FlexGroup
SnapMirror relationships.
Is Adaptive Enabled: Whether or not adaptive is enabled.
Can be:
- true: Adaptive is enabled.
- false: Adaptive is disabled.
This parameter is supported only for FlexVol
SnapMirror relationships between ONTAP
endpoints.
Backoff level: Level of backoff of SnapMirror relationship
transfers in the presence of client ops.
Can be:
- high
- medium
- none
This parameter is supported only for FlexVol
SnapMirror relationships.

Number of Successful Updates: The number of successful SnapMirror update
operations for the relationship since the
last reboot, aggregate relocation, takeover/
giveback, or metrocluster switchover/switchback
involving the node that hosts the relationship.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Number of Failed Updates: The number of failed SnapMirror update
operations for the relationship since the
last reboot, aggregate relocation, takeover/
giveback, or metrocluster switchover/switchback
involving the node that hosts the relationship.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Number of Successful Resyncs: The number of successful SnapMirror resync
operations for the relationship since the
last reboot, aggregate relocation, takeover/
giveback, or metrocluster switchover/switchback
involving the node that hosts the relationship.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Number of Failed Resyncs: The number of failed SnapMirror resync
operations for the relationship since the
last reboot, aggregate relocation, takeover/
giveback, or metrocluster switchover/switchback
involving the node that hosts the relationship.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Number of Successful Breaks: The number of successful SnapMirror break
operations for the relationship since the
last reboot, aggregate relocation, takeover/
giveback, or metrocluster switchover/switchback
involving the node that hosts the relationship.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Number of Failed Breaks: The number of failed SnapMirror break
operations for the relationship since the
last reboot, aggregate relocation, takeover/
giveback, or metrocluster switchover/switchback

involving the node that hosts the relationship.
Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Total Transfer Bytes: Cumulative bytes transferred for the
relationship. Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.
Total Transfer Time: Cumulative total transfer time in seconds for
the relationship since the last reboot, aggregate
relocation, takeover/giveback, or metrocluster
switchover/switchback involving the node that
hosts the relationship. Only for relationships with
"Relationship Capability"
of "8.2 and above" .
This parameter is not supported for Vserver
SnapMirror relationships.

```
== 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.

{ ``[-S, -source-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Source Path::
Select SnapMirror relationships that have a matching source path name.

| ``[-source-cluster <Cluster name>]`` - Source Cluster::
Select SnapMirror relationships that have a matching source cluster name.
This parameter is not supported for relationships with non-ONTAP,
FlexGroup volume, or Vserver source endpoints.

``[-source-vserver <vserver name>]`` - Source Vserver::
Select SnapMirror relationships that have a matching source Vserver name.
This parameter is not supported for relationships with non-ONTAP source
endpoints.

``[-source-volume <volume name>]`` - Source Volume }::
```

```
Select SnapMirror relationships that have a matching source volume name.
This parameter is not supported for relationships with non-ONTAP source
endpoints.

{ ``[-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Destination
Path::
Select SnapMirror relationships that have a matching destination path
name.
+
NOTE: Using wildcards with this parameter:
+
  * To match all Vserver SnapMirror relationships, use: `-destination-
path```_*:_``
  * To match all the SnapMirror relationships except Vserver SnapMirror
relationships in the cluster, use: `-destination-path```_*_``



+



| ``[-destination-cluster <Cluster name>]`` - Destination Cluster::
Select SnapMirror relationships that have a matching destination cluster
name. This parameter is not supported for relationships with non-ONTAP,
FlexGroup volume, or Vserver destination endpoints.

``[-destination-vserver <vserver name>]`` - Destination Vserver::
Select SnapMirror relationships that have a matching destination Vserver
name. This parameter is not supported for relationships with non-ONTAP
destination endpoints.

``[-destination-volume <volume name>]`` - Destination Volume }::
Select SnapMirror relationships that have a matching destination volume
name. This parameter is not supported for relationships with non-ONTAP
destination endpoints.

``[-type <snapmirrorType>]`` - Relationship Type::
Select SnapMirror relationships that have a matching relationship type.
Vservers support only DP SnapMirror relationships. Possible values are:
+
  * DP
  * LS
  * XDP
  * TDP
  * RST
```

```
``[-relationship-group-type
{none|vserver|consistencygroup|flexgroup|vserverflexgroup|vserverconsisten
cygroup}]`` - Relationship Group Type::
Select SnapMirror relationships that have a matching relationship group
type. Possible values are:
+
  * ``_none_``
  * ``_vserver_``
  * ``_flexgroup_``
  * ``_consistencygroup_``



``[-vserver <vserver name>]`` - Managing Vserver::
Select SnapMirror relationships that have a matching managing Vserver
name. The `-vserver` option is currently a reserved option.

``[-schedule <text>]`` - SnapMirror Schedule::
Select SnapMirror relationships that have a matching schedule.

``[-policy-type {vault|async-mirror|mirror-vault|strict-sync-mirror|sync-
mirror|automated-failover|automated-failover-duplex|continuous}]`` -
SnapMirror Policy Type::
Selects SnapMirror relationships that have a matching SnapMirror policy
type. Possible values are:
+
  * async-mirror
  * vault
  * mirror-vault
  * automated-failover
  * automated-failover-duplex
  * sync-mirror
  * strict-sync-mirror



``[-policy <sm_policy>]`` - SnapMirror Policy::
Select SnapMirror relationships that have a matching SnapMirror policy.

``[-tries <integer_or_unlimited>]`` - Tries Limit::
Select SnapMirror relationships that have a matching tries limit.

``[-k, -throttle <throttleType>]`` - Throttle (KB/sec)::
Select SnapMirror relationships that have a matching throttle.
```

```
``[-cg-item-mappings {<source volume>:@<destination volume>|<source item
path>:@<destination item path>}]`` - Consistency Group Item Mappings::
Select SnapMirror relationships that have a matching cg-item-mappings.

``[-current-throttle <throttleType>]`` - Current Transfer Throttle
(KB/sec)::
Select SnapMirror relationships that have a matching current throttle.

``[-state <mirror state>]`` - Mirror State::
Select SnapMirror relationships that have a matching mirror state.
Possible values are:
+
  * Uninitialized
  * Snapmirrored
  * Broken-off



``[-status <mirror status>]`` - Relationship Status::
Select SnapMirror relationships that have a matching relationship status.
Possible values are:
+
  * Idle
  * Queued
  * Transferring
  * Preparing
  * Finalizing
  * Aborting
  * Quiesced
  * Quiescing
  * Checking

+
Status values Finalizing, Checking, Waiting and Preparing are not
supported for Vserver SnapMirror relationships.

``[-file-restore-file-count <integer>]`` - File Restore File Count::
The number of files being restored by file restore.

``[-file-restore-file-list <text>,...]`` - File Restore File List::
List of the destination file names of the files being restored by file
restore.

``[-transfer-snapshot <text>]`` - Transfer Snapshot::
Select SnapMirror relationships that have a matching transfer snapshot.

``[-snapshot-progress {<integer>[KB|MB|GB|TB|PB]}]`` - Snapshot Progress::
```

```
Select SnapMirror relationships that have a matching snapshot progress.

``[-total-progress {<integer>[KB|MB|GB|TB|PB]}]`` - Total Progress::
Select SnapMirror relationships that have a matching total progress.

``[-network-compression-ratio <text>]`` - Network Compression Ratio::
Select SnapMirror relationships that have a matching network compression
ratio. This parameter is not supported for Vserver SnapMirror
relationships.

``[-snapshot-checkpoint {<integer>[KB|MB|GB|TB|PB]}]`` - Snapshot
Checkpoint::
Select SnapMirror relationships that have a matching snapshot checkpoint.
This parameter is not supported for Vserver SnapMirror relationships.

``[-newest-snapshot <text>]`` - Newest Snapshot::
Select SnapMirror relationships that have a matching newest snapshot.

``[-newest-snapshot-timestamp <MM/DD HH:MM:SS>]`` - Newest Snapshot
Timestamp::
Select SnapMirror relationships that have a matching newest snapshot
timestamp.

``[-exported-snapshot <text>]`` - Exported Snapshot::
Select SnapMirror relationships that have a matching exported snapshot
name. For load-sharing mirror relationships, if the exported-snapshot
field for a relationship has a dash (-), the load-sharing mirror is
lagging behind the up-to-date mirrors in the set.

``[-exported-snapshot-timestamp <MM/DD HH:MM:SS>]`` - Exported Snapshot
Timestamp::
Select SnapMirror relationships that have a matching exported snapshot
timestamp.

``[-healthy {true|false}]`` - Healthy::
Select SnapMirror relationships that have a matching healthy condition.

``[-relationship-id <UUID>]`` - Relationship ID::
Select SnapMirror relationships that have a matching relationship ID. This
parameter is not supported for Vserver SnapMirror relationships.

``[-current-operation-id <UUID>]`` - Current Operation ID::
Select SnapMirror relationships that have a matching operation unique
identifier of the currently executing SnapMirror operation.

``[-current-transfer-type
{initialize|update|resync|restore|check|file_restore|cggrs_initialize|cggr
```

```
s_resync|cg_update|cg_initialize|cg_resync|cg_restore|catalog_metadata|cg_
file_restore|failover|aufo|failover_incapable|cg_migrate}]`` - Transfer
Type::
Select SnapMirror relationships that have a matching current transfer
type.

``[-current-transfer-error <text>]`` - Transfer Error::
Select SnapMirror relationships that have a matching current transfer
error.

``[-last-transfer-type
{initialize|update|resync|restore|check|file_restore|cggrs_initialize|cggr
s_resync|cg_update|cg_initialize|cg_resync|cg_restore|catalog_metadata|cg_
file_restore|failover|aufo|failover_incapable|cg_migrate}]`` - Last
Transfer Type::
Select SnapMirror relationships that have a matching last transfer type.

``[-last-transfer-error <text>]`` - Last Transfer Error::
Select SnapMirror relationships that have a matching last transfer error.

``[-last-transfer-size {<integer>[KB|MB|GB|TB|PB]}]`` - Last Transfer
Size::
Select SnapMirror relationships that have a matching last transfer size.

``[-last-transfer-network-compression-ratio <text>]`` - Last Transfer
Network Compression Ratio::
Select SnapMirror relationships that have a matching last transfer network
compression ratio. This parameter is not supported for Vserver SnapMirror
relationships.

``[-last-transfer-duration <[[<hours>:]<minutes>:]<seconds>>]`` - Last
Transfer Duration::
Select SnapMirror relationships that have a matching last transfer
duration.

``[-last-transfer-from <text>]`` - Last Transfer From::
Select SnapMirror relationships that have a matching last transfer source.

``[-last-transfer-end-timestamp <MM/DD HH:MM:SS>]`` - Last Transfer End
Timestamp::
Select SnapMirror relationships that have a matching last transfer end
timestamp.

``[-unhealthy-reason <text>]`` - Unhealthy Reason::
Select SnapMirror relationships that have a matching unhealthy reason.

``[-progress-last-updated <MM/DD HH:MM:SS>]`` - Progress Last Updated::
```

```
Select SnapMirror relationships that have a matching progress last
updated.

``[-relationship-capability <text>]`` - Relationship Capability::
Select SnapMirror relationships that have a matching relationship
capability. This parameter is not supported for Vserver SnapMirror
relationships.

``[-lag-time <[[<hours>:]<minutes>:]<seconds>>]`` - Lag Time::
Select SnapMirror relationships that have a matching lag time.

``[-current-transfer-priority {low|normal}]`` - Current Transfer
Priority::
Select SnapMirror relationships that have a matching current transfer
priority.

``[-is-smtape-op {true|false}]`` - SMTape Operation::
Select SnapMirror relationships that have a matching smtape operation.

``[-destination-volume-node <nodename>]`` - Destination Volume Node Name::
Select SnapMirror relationships that have a matching destination volume
node name. This parameter is not supported for Vserver SnapMirror
relationships.

``[-identity-preserve {true|false}]`` - Identity Preserve Vserver DR::
Select SnapMirror relationships that have a matching value for identity-
preserve. This parameter is valid only for Vserver SnapMirror
relationships.

``[-expand <true>]`` - Show Constituents of the Group::
Specifies whether to display constituent relationships of Vserver,
FlexGroup and Consistency Group SnapMirror relationships. By default, the
constituents are not displayed.

``[-update-successful-count <integer>]`` - Number of Successful Updates::
Select SnapMirror relationships that have a matching number of successful
updates. This parameter is not supported for Vserver SnapMirror
relationships.

``[-update-failed-count <integer>]`` - Number of Failed Updates::
Select SnapMirror relationships that have a matching number of failed
updates. This parameter is not supported for Vserver SnapMirror
relationships.

``[-resync-successful-count <integer>]`` - Number of Successful Resyncs::
Select SnapMirror relationships that have a matching number of successful
resyncs. This parameter is not supported for Vserver SnapMirror
```

```
relationships.

``[-resync-failed-count <integer>]`` - Number of Failed Resyncs::
Select SnapMirror relationships that have a matching number of failed
resyncs. This parameter is not supported for Vserver SnapMirror
relationships.

``[-break-successful-count <integer>]`` - Number of Successful Breaks::
Select SnapMirror relationships that have a matching number of successful
breaks. This parameter is not supported for Vserver SnapMirror
relationships.

``[-break-failed-count <integer>]`` - Number of Failed Breaks::
Select SnapMirror relationships that have a matching number of failed
breaks. This parameter is not supported for Vserver SnapMirror
relationships.

``[-total-transfer-bytes <integer>]`` - Total Transfer Bytes::
Select SnapMirror relationships that have a matching total transfer bytes.
This parameter is not supported for Vserver SnapMirror relationships.

``[-total-transfer-time-secs <integer>]`` - Total Transfer Time in
Seconds::
Select SnapMirror relationships that have a matching total transfer time
in seconds. This parameter is not supported for Vserver SnapMirror
relationships.

``[-msid-preserve {true|false}]`` - Source Volume MSIDs Preserved::
This parameter specifies whether the volume MSIDs are preserved at the
destination. This parameter is applicable only for Vserver SnapMirror
relationships.

``[-is-auto-expand-enabled {true|false}]`` - Is Auto Expand Enabled::
Select SnapMirror relationships that have a matching value for auto
expand. This parameter is supported only for FlexGroup SnapMirror
relationships. Possible values are:
+
  * true
  * false



``[-percent-complete-cur-status <integer>]`` - Percent Complete for
Current Status::
Select SnapMirror relationships that have a matching percent complete for
the current status.
```

```
``[-backoff-level {high|medium|none}]`` - SM Backoff Level due to Client
Ops::
Select SnapMirror relationships that have a matching value for backoff
level. This parameter is supported only for FlexVol SnapMirror
relationships. Possible values are:
+
  * high
  * medium
  * none



== Examples

The example below displays summary information for all SnapMirror
relationships with destination endpoints in the current cluster:
```

cluster2::> snapmirror show
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
---------- ---- ---------- ------- ---------- --------- ------- --------
cluster1-vs2.example1.com:
DP cluster2-dvs2.example2.com:
Snapmirrored
Idle - true -
cluster2-vs1.example.com:dp_src1
DP cluster2-vs2.example.com:dp_dst1
Snapmirrored
Idle - true -
cluster2-vs1.example.com:xdp_src1
XDP cluster2-vs2.example.com:xdp_dst1
Snapmirrored
Idle - true -
cluster2://cluster2-vs1.example.com/ls_src1
LS cluster2://cluster2-vs1.example.com/ls_mr1
Snapmirrored
Idle - true -
cluster2://cluster2-vs1.example.com/ls_mr2
Snapmirrored
Idle - true -
5 entries were displayed.

```
The example below displays detailed information for the SnapMirror
relationship with the destination endpoint ``_cluster2-
vs2.example.com:dp_dst1_`` .
```

cluster2::> snapmirror show -destination-path cluster2-vs2.example.com:dp_dst1
Source Path: cluster2-vs1.example.com:dp_src1
Destination Path: cluster2-vs2.example.com:dp_dst1

Relationship Type: DP
Relationship Group Type: none
SnapMirror Schedule: -
SnapMirror Policy Type: async-mirror
SnapMirror Policy: DPDefault
Tries Limit: -
Throttle (KB/sec): unlimited
Mirror State: Snapmirrored
Relationship Status: Idle
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Percent Complete for Current Status: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: snapmirror.3d19af37-8f5e-11e1-8c83-123478563412_2147484676.2012-04-27_025137
Newest Snapshot Timestamp: 04/27 02:51:42
Exported Snapshot: snapmirror.3d19af37-8f5e-11e1-8c83-123478563412_2147484676.2012-04-27_025137
Exported Snapshot Timestamp: 04/27 02:51:42
Healthy: true
Unhealthy Reason: -
Destination Volume Node: cluster2-node1
Relationship ID: cdc70a81-8f5f-11e1-8392-123478563412
Current Operation ID: -
Transfer Type: -
Transfer Error: -
Current Throttle: -
Current Transfer Priority: -
Last Transfer Type: update
Last Transfer Error: -
Last Transfer Size: 530.2MB
Last Transfer Network Compression Ratio: 111.7:1
Last Transfer Duration: 0:2:53
Last Transfer From: cluster2-vs1.example.com:dp_src1
Last Transfer End Timestamp: 04/27 02:51:45
Progress Last Updated: -
Relationship Capability: 8.2 and above
Lag Time: 133:50:40
Identity Preserve Vserver DR: -
Volume MSIDs Preserved: -
Is Auto Expand Enabled: -
Is Adaptive: -
Number of Successful Updates: 1
Number of Failed Updates: 0
Number of Successful Resyncs: 0
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
Total Transfer Bytes: 663552
Total Transfer Time in Seconds: 3

```
The example below displays detailed information for SnapMirror
relationships with the ``_Relationship Capability_`` of ``_"Pre 8.2"_``
source or destination endpoints in the current cluster.
```

cluster2::> snapmirror show -relationship-capability "Pre 8.2" -instance
Source Path: cluster2://cluster2-vs1.example.com/ls_src1
Destination Path: cluster2://cluster2-vs1.example.com/ls_mr1
Relationship Type: LS
Relationship Group Type: -
SnapMirror Schedule: -
SnapMirror Policy Type: -
SnapMirror Policy: -
Tries Limit: 8
Throttle (KB/sec): unlimited
Mirror State: Snapmirrored
Relationship Status: Idle
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Percent Complete for Current Status: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: snapmirror.3d4e52c5-8f5c-11e1-8392-123478563412_3_2147484684.2012-05-02_163506
Newest Snapshot Timestamp: 05/02 16:35:06
Exported Snapshot: snapmirror.3d4e52c5-8f5c-11e1-8392-123478563412_3_2147484684.2012-05-02_163506
Exported Snapshot Timestamp: 05/02 16:35:06
Healthy: true
Unhealthy Reason: -
Destination Volume Node: -
Relationship ID: -
Current Operation ID: -
Transfer Type: -
Transfer Error: -
Last Transfer Type: -
Last Transfer Error: -
Last Transfer Size: -
Last Transfer Network Compression Ratio: -
Last Transfer Duration: -
Last Transfer From: -
Last Transfer End Timestamp: -
Progress Last Updated: -
Relationship Capability: Pre 8.2
Lag Time: -
SnapMirror Policy: -
Identity Preserve Vserver DR: -
Volume MSIDs Preserved: -
Is Auto Expand Enabled: -
Is Adaptive: -
Number of Successful Updates: -
Number of Failed Updates: -
Number of Successful Resyncs: -

Number of Failed Resyncs: -
Number of Successful Breaks: -
Number of Failed Breaks: -
Total Transfer Bytes: -
Total Transfer Time in Seconds: -
Source Path: cluster2://cluster2-vs1.example.com/ls_src1
Destination Path: cluster2://cluster2-vs1.example.com/ls_mr2
Relationship Type: LS
Relationship Group Type: -
SnapMirror Schedule: -
SnapMirror Policy Type: -
SnapMirror Policy: -
Tries Limit: 8
Throttle (KB/sec): unlimited
Mirror State: Snapmirrored
Relationship Status: Idle
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Percent Complete for Current Status: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: snapmirror.3d4e52c5-8f5c-11e1-8392-123478563412_3_2147484684.2012-05-02_163506
Newest Snapshot Timestamp: 05/02 16:35:06
Exported Snapshot: snapmirror.3d4e52c5-8f5c-11e1-8392-123478563412_3_2147484684.2012-05-02_163506
Exported Snapshot Timestamp: 05/02 16:35:06
Healthy: true
Unhealthy Reason: -
Destination Volume Node: -
Relationship ID: -
Current Operation ID: -
Transfer Type: -
Transfer Error: -
Last Transfer Type: -
Last Transfer Error: -
Last Transfer Size: -
Last Transfer Network Compression Ratio: -
Last Transfer Duration: -
Last Transfer From: -
Last Transfer End Timestamp: -
Progress Last Updated: -
Relationship Capability: Pre 8.2
Lag Time: -
SnapMirror Policy: -
Identity Preserve Vserver DR: -
Volume MSIDs Preserved: -
Is Auto Expand Enabled: -
Is Adaptive: -
Number of Successful Updates: -
Number of Failed Updates: -
Number of Successful Resyncs: -
Number of Failed Resyncs: -
Number of Successful Breaks: -

Number of Failed Breaks: -
Total Transfer Bytes: -
Total Transfer Time in Seconds: -

2 entries were displayed.

```
The example below displays detailed information for the Vserver SnapMirror
relationship with the destination endpoint ``_cluster2-
dvs2.example2.com:_`` .
```

cluster2::> snapmirror show -destination-path cluster2-dvs2.example2.com:
Source Path: cluster1-vs2.example1.com:
Destination Path: cluster2-dvs2.example2.com:
Relationship Type: DP
Relationship Group Type: -
SnapMirror Schedule: -
SnapMirror Policy Type: async-mirror
SnapMirror Policy: DPDefault
Tries Limit: -
Throttle (KB/sec): unlimited
Mirror State: Snapmirrored
Relationship Status: Idle
File Restore File Count: -
File Restore File List: -
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Percent Complete for Current Status: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: vserverdr.1d519e9c-7838-11e3-91fb-123478563412.2014-01-13_110707
Newest Snapshot Timestamp: 01/13 11:07:07
Exported Snapshot: vserverdr.1d519e9c-7838-11e3-91fb-123478563412.2014-01-13_110707
Exported Snapshot Timestamp: 01/13 11:07:07
Healthy: true
Unhealthy Reason: -
Destination Volume Node: -
Relationship ID: -
Operation ID: -
Transfer Type: -
Transfer Error: -
Current Throttle: -
Current Transfer Priority: -
Last Transfer Type: resync
Last Transfer Error: -
Last Transfer Size: -
Last Transfer Network Compression Ratio: -
Last Transfer Duration: -
Last Transfer From: cluster1-vs2.example1.com:
Last Transfer End Timestamp: -
Progress Last Updated: -
Relationship Capability: -

Lag Time: 18:47:9
Identity Preserve Vserver DR: false
Volume MSIDs Preserved: true
Is Auto Expand Enabled: -
Is Adaptive: -
Number of Successful Updates: -
Number of Failed Updates: -
Number of Successful Resyncs: -
Number of Failed Resyncs: -
Number of Successful Breaks: -
Number of Failed Breaks: -
Total Transfer Bytes: -
Total Transfer Time in Seconds: -

```
The following example displays detailed information for the SnapMirror
relationship with the AltaVault destination endpoint
``_10.0.0.11:/share/share1_`` :
```

cluster2::> snapmirror show -destination-path 10.0.0.11:/share/share1
Source Path: cluster2-vs1.example.com:data_ontap_vol
Destination Path: 10.0.0.11:/share/share1
Relationship Type: XDP
Relationship Group Type: none
SnapMirror Schedule: -
SnapMirror Policy Type: vault
SnapMirror Policy: XDPDefault
Tries Limit: -
Throttle (KB/sec): unlimited
Mirror State: Snapmirrored
Relationship Status: Idle
Transfer Snapshot: -
Snapshot Progress: -
Total Progress: -
Percent Complete for Current Status: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: snapmirror.3d19af37-8f5e-11e1-8c83-123478563412_2147484676.2012-04-27_025137
Newest Snapshot Timestamp: 04/27 02:51:42
Exported Snapshot: snapmirror.3d19af37-8f5e-11e1-8c83-123478563412_2147484676.2012-04-27_025137
Exported Snapshot Timestamp: 04/27 02:51:42
Healthy: true
Unhealthy Reason: -
Destination Volume Node: -
Relationship ID: cdc70a81-8f5f-11e1-8392-123478563463
Current Operation ID: -
Transfer Type: -
Transfer Error: -
Current Throttle: -
Current Transfer Priority: -
Last Transfer Type: update
Last Transfer Error: -
Last Transfer Size: 530.2MB

Last Transfer Network Compression Ratio: 1:1
Last Transfer Duration: 0:2:53
Last Transfer From: cluster2-vs1.example.com:data_ontap_vol
Last Transfer End Timestamp: 04/27 02:51:45
Progress Last Updated: -
Relationship Capability: 8.2 and above
Lag Time: 133:50:40
Identity Preserve Vserver DR: -
Volume MSIDs Preserved: -
Is Auto Expand Enabled: -
Is Adaptive: -
Number of Successful Updates: 1
Number of Failed Updates: 0
Number of Successful Resyncs: 0
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
Total Transfer Bytes: 663552
Total Transfer Time in Seconds: 3

```
The example shows the usage of the ``_-expand_`` parameter to additionally
display the constituents of Vserver SnapMirror relationships with
destination endpoints in the current cluster. Note that in the following
example, since there is no volume level relationship for the root volume
of a Vserver, it is not shown in the output:
```

cluster2::> snapmirror show -expand
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
----------- ---- ------------ ------- -------------- --------- ------- --------
cluster1-vs1.example1.com:
DP cluster2-dvs1.example2.com:
Snapmirrored
Idle - true -
cluster1-vs1.example1.com:vol1
DP cluster2-dvs1.example2.com:vol1
Snapmirrored
Idle - true -
cluster1-vs2.example1.com:
DP cluster2-dvs2.example2.com:
Snapmirrored
Idle - true -
cluster1-vs2.example1.com:vol1
DP cluster2-dvs2.example2.com:vol1
Snapmirrored
Idle - true -
4 entries were displayed.

```
The example below displays summary information for SnapMirror
relationships with Consistency Group destination endpoints in the current
cluster:
```

cluster2::> snapmirror show
Progress
Source Destination Mirror Relationship Total Last
Path Type Path State Status Progress Healthy Updated
---------- ---- ---------- ------- ---------- --------- ------- --------
vs1.example.com:/cg/cg1
XDP vs2.example.com:/cg/cg1_dst
Snapmirrored
Idle - true -

```
The following example displays detailed information for the SnapMirror
relationship with destination Consistency Group ``_cg1_dst_`` in Vserver
``_vs2.example.com_`` :
```

cluster2::> snapmirror show -destination-path vs2.example.com:/cg/cg1_dst
Source Path: vs1.example.com:/cg/cg1
Source Cluster: -
Source Vserver: vs1.example.com
Source Volume: -
Mirror Type: XDP
Destination Cluster ID: -
Destination Vserver ID: -
Destination Vserver UUID: 69836aef-7aef-11ef-9229-005056ac6055
Destination Volume MSID: -
Destination Volume DSID: -
Destination Path: vs2.example.com:/cg/cg1_dst
Destination Cluster: -
Destination Vserver: vs2.example.com
Destination Volume: -
Relationship Type: XDP
Relationship Group Type: consistencygroup
Managing Vserver: vs2.example.com
SnapMirror Schedule: -
SnapMirror Policy Type: mirror-vault
SnapMirror Policy: MirrorAndVault
Tries Limit: -
Throttle (KB/sec): unlimited
Consistency Group Item Mappings: cg1_lun1:@cg1_lun1_dst,
cg1_lun2:@cg1_lun2_dst
Current Transfer Throttle (KB/sec): -
Mirror State: Snapmirrored
Relationship Status: Idle
File Restore File Count: -
File Restore File List: -
Transfer Snapshot: -
Snapshot Progress: -

Total Progress: -
Network Compression Ratio: -
Snapshot Checkpoint: -
Newest Snapshot: snapmirrorCG.69836aef-7aef-11ef-9229-005056ac6055_2149683835.2024-09-30_083931
Newest Snapshot Timestamp: 09/30 08:40:39
Exported Snapshot: snapmirrorCG.69836aef-7aef-11ef-9229-005056ac6055_2149683835.2024-09-30_083931
Exported Snapshot Timestamp: 09/30 08:40:39
Healthy: true
Physical Replica: -
Relationship ID: 094b1955-7f29-11ef-8894-005056ac6055
Source Vserver UUID: 6b6bf517-7aef-11ef-b7ab-005056ac77f1
Destination Vserver UUID: 69836aef-7aef-11ef-9229-005056ac6055
Current Operation ID: -
Transfer Type: -
Transfer Error: -
Last Transfer Type: initialize
Last Transfer Error: -
Last Transfer Error Codes: -
Last Transfer Size: 74.38KB
Last Transfer Network Compression Ratio: 1:1
Last Transfer Duration: 0:1:14
Last Transfer From: vs1.example.com:/cg/cg1
Last Transfer End Timestamp: 09/30 08:40:45
Unhealthy Reason: -
Progress Last Updated: -
Relationship Capability: 8.2 and above
Lag Time: 0:0:58
Current Transfer Priority: -
SMTape Operation: -
Constituent Relationship: false
Destination Volume Node Name: -
Identity Preserve Vserver DR: -
Number of Successful Updates: 0
Number of Failed Updates: 0
Number of Successful Resyncs: 0
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
Total Transfer Bytes: 76160
Total Transfer Time in Seconds: 74
Source Volume MSIDs Preserved: -
OpMask: ffffffffffffffff
Is Auto Expand Enabled: true
FabricLink Source Bucket: -
FabricLink Source Role: -
FabricLink Destination Bucket: -
FabricLink Destination Role: -
FabricLink Topology UUID: -
FabricLink Pull Byte Count: -
FabricLink Push Byte Count: -
FabricLink Pending Work Task Count: -
FabricLink Status: -
Percent Complete for Current Status: -

SM Backoff Level due to Client Ops: -

```
== Related Links

* xref:{relative_path}snapmirror-list-destinations.html[snapmirror list-
destinations]
* xref:{relative_path}snapmirror-policy-create.html[snapmirror policy
create]




[[ID87cf6ff2e7d536a059c090428e38ae9c]]
= snapmirror update-ls-set
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start an incremental load-sharing set transfer

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror update-ls-set` command updates a set of load-sharing
mirrors. The command makes destination volumes, in the group of load-
sharing mirrors, up-to-date mirrors of the source volume.

The key parameter that identifies the set of load-sharing mirrors is the
source volume. SnapMirror transfers are performed from the source volume
to each of the up-to-date destination volumes in the set of load-sharing
mirrors.

The `snapmirror update-ls-set` command performs an incremental transfer to
each of the destination volumes. During an incremental transfer, ONTAP
takes a snapshot on the source volume to capture the current image of the
source volume, finds the most recent common snapshot between the source
and destination volumes, and incrementally transfers snapshots that are
newer than the common snapshot to the destination volume.

NOTE: You still need to use the `snapmirror update-ls-set` command to
manually update the set of load-sharing mirrors even if the set only has
one destination mirror. The xref:{relative_path}snapmirror-
```

```
update.html[snapmirror update] command can only be used to bring up to
date a specific destination mirror that is lagging to the set.



After an update using the `snapmirror update-ls-set` command successfully
completes, the last snapshot transferred is made the new exported snapshot
on the destination volumes.

This command is only supported for SnapMirror relationships with the field
``_"Relationship Capability"_`` showing as ``_"Pre 8.2"_`` in the output
of the xref:{relative_path}snapmirror-show.html[snapmirror show] command.

== Parameters

{ ``-S, -source-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}`` - Source Path::
This parameter specifies the source endpoint of the SnapMirror
relationship in one of four path formats. The normal format includes the
names of the Vserver (vserver) and/or volume (volume). To support
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , a
format which also includes the name of the cluster (cluster) is provided.
The ``_"Pre 8.2"_`` format cannot be used when operating in a Vserver
context on relationships with ``_"Relationship Capability"_`` of ``_"8.2
and above"_`` . For SnapMirror relationships with an AltaVault source, the
source endpoint is specified in the form ``_hostip:/share/share-name_`` .
For SnapMirror relationships with a SolidFire source, the source endpoint
is specified in the form ``_hostip:/lun/name_`` .

| ``[-source-cluster <Cluster name>]`` - Source Cluster::
Specifies the source cluster of the SnapMirror relationship. If this
parameter is specified, the `-source-vserver` and `-source-volume`
parameters must also be specified. This parameter is only applicable for
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .
This parameter cannot be specified when operating in a Vserver context on
relationships with ``_"Relationship Capability"_`` of ``_"8.2 and
above"_`` .

``-source-vserver <vserver name>`` - Source Vserver::
Specifies the source Vserver of the SnapMirror relationship. For
relationships with volumes as endpoints, if this parameter is specified,
parameters `-source-volume` and for relationships with ``_"Relationship
Capability"_`` of ``_"Pre 8.2"_`` , `-source-cluster` must also be
specified. This parameter is not supported for relationships with non-
ONTAP source endpoints.
```

```
``-source-volume <volume name>`` - Source Volume }::
Specifies the source volume of the SnapMirror relationship. If this
parameter is specified, parameters `-source-vserver` and for relationships
with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-source-
cluster` must also be specified. This parameter is not supported for
relationships with non-ONTAP source endpoints.

``[-w, -foreground <true>]`` - Foreground Process::
This specifies whether the operation runs as a foreground process. If this
parameter is specified, the default setting is ``_true_`` (the operation
runs in the foreground). When set to ``_true_`` , the command will not
return until the process completes. This parameter is only applicable to
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .

== Examples

To update the group of load-sharing mirrors for the source endpoint named
``_//vs1.example.com/vs1_root_`` , type the following command:
```

cluster1::> snapmirror update-ls-set -source-path //vs1.example.com/vs1_root

```
== Related Links

* xref:{relative_path}snapmirror-update.html[snapmirror update]
* xref:{relative_path}snapmirror-show.html[snapmirror show]




[[ID6aa3fcb54b6036a61d4d927d93ccf16c]]
= snapmirror update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start an incremental transfer

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror update` command updates the destination volume or non-Data
ONTAP endpoint of a SnapMirror relationship. The `snapmirror update`
```

```
command behaves differently for data protection (DP), extended data
protection (XDP) and load-sharing (LS) relationships. Refer to the `-type`
parameter of the xref:{relative_path}snapmirror-create.html[snapmirror
create] command to understand different types of relationships supported
by SnapMirror.

The `snapmirror update` command performs an incremental transfer.

Before using this command, the relationship must be initialized using the
xref:{relative_path}snapmirror-initialize.html[snapmirror initialize] or
xref:{relative_path}snapmirror-initialize-ls-set.html[snapmirror
initialize-ls-set] commands.

For data protection SnapMirror relationships with volumes as endpoints,
the `snapmirror update` command makes the destination volume an up-to-date
mirror of the source volume with the following steps:

* If the source volume is read-write, takes a snapshot on the source
volume to capture the current image of the source volume
* Finds the most recent snapshot on the destination volume and validates
that the corresponding snapshot is still present on the source
* Incrementally transfers snapshots that are newer than the corresponding
snapshot to the destination volume



You can use the `snapmirror update` command to update a specific load-
sharing mirror that lags behind up-to-date destination volumes in the set
of load-sharing mirrors. An update to the lagging load-sharing mirror
should bring it up to date with the other up-to-date destination volumes
in the set of load-sharing mirrors.

NOTE: Using the `snapmirror update` command to update a set of load-
sharing mirrors will not work. Use the xref:{relative_path}snapmirror-
update-ls-set.html[snapmirror update-ls-set] command to update a set of
load-sharing mirrors.



For extended data protection (XDP) relationships with a `snapmirror
policy` of type ``_async-mirror_`` , a `snapmirror update` always creates
a new snapshot on the source volume. Depending on the rules in the policy,
the command might transfer just the newly created snapshot or all
snapshots that are newer than the common snapshot including the newly
created snapshot to the destination volume.

For extended data protection (XDP) relationships with a `snapmirror
```

policy` of type ``_vault_`` , a `snapmirror update` does not create a new snapshot on the source volume but transfers only selected snapshots that are newer than the common snapshot to the destination volume. (Those older than the common copy can be transferred by using the `-source-snapshot` parameter.) snapshots are selected by matching the value of `-snapmirror-label` of a snapshot with the value of `-snapmirror-label` of one of the rules from the corresponding SnapMirror policy associated with the SnapMirror relationship. All matching snapshots are incrementally transferred to the destination volume.

For extended data protection (XDP) relationships with a `snapmirror policy` of type ``_mirror-vault_`` , a `snapmirror update` always creates a new snapshot on the source volume and transfers only selected snapshots that are newer than the common snapshot. The newly created snapshot is always selected.

For extended data protection (XDP) relationships with a `snapmirror policy` of type ``_vault_`` or ``_mirror-vault_`` , the `snapmirror update` command also manages expiration of snapshots on the destination volume. It does so by deleting snapshots that have exceeded the value of `-keep` for the matching rule from the corresponding SnapMirror policy associated with the SnapMirror relationship. snapshots that match the same `-snapmirror-label` will be deleted in oldest-first order.

For relationships with a policy of type ``_strict-sync-mirror_`` or ``_sync-mirror_`` , this command creates a new common snapshot and designates it as the exported snapshot on the destination volume. It updates the destination read-only view because IO is redirected to the new exported snapshot. Clients could experience a brief latency spike during this process as the primary IO is temporarily fenced. This command is allowed only when the relationship-status is InSync. The command retains two pairs of common snapshots and deletes the older ones.

For SnapMirror active sync relationships with Consistency Groups of FlexVol volumes as endpoints, the `snapmirror update` command creates a new coordinated common snapshot. The command retains two coordinated common snapshots and deletes the oldest one.

For data protection relationships, the parameter `-source-snapshot` is optional and only allows for the transfer of snapshots newer than the common snapshot up to the specified `-source-snapshot` .

For extended data protection (XDP) relationships the parameter `-source-snapshot` is optional.

For extended data protection (XDP) relationships with a `snapmirror

```
policy` of type ``_vault_`` or ``_mirror-vault_`` , the parameter `-
source-snapshot` allows transfer of a snapshot that is older than the
common snapshot and/or might not be selected for transfer based on policy-
based selection of a scheduled update transfer.

For extended data protection (XDP) relationships with a `snapmirror
policy` of type ``_async-mirror_`` , the `snapmirror update` with
parameter `-source-snapshot` does not create a new snapshot on the source
volume. Depending on the rules in the policy, the command might transfer
just the specified snapshot or snapshots that are newer than the common
snapshot up to and including the specified snapshot to the destination
volume.

After the `snapmirror update` command successfully completes, the last
snapshot transferred is designated as the new exported snapshot on the
destination volume. If an update to an extended data protection (XDP)
relationship specifies a snapshot using the `-source-snapshot` parameter
that is older than the common snapshot, after the `snapmirror update`
successfully completes, the exported snapshot on the destination volume
will remain unchanged.

If the `snapmirror update` does not finish successfully--for example, due
to a network failure or because a xref:{relative_path}snapmirror-
abort.html[snapmirror abort] command was issued--a restart checkpoint
might be recorded on the destination volume. If a restart checkpoint is
recorded, the next update restarts and continues the transfer from the
restart checkpoint. For extended data protection (XDP) relationships, the
next update will restart and continue the old transfer regardless of
whether the snapshot being transferred is a matching snapshot or not.

This command is supported for SnapMirror relationships with the field
``_"Relationship Capability"_`` showing as either ``_"8.2 and above"_`` or
``_"Pre 8.2"_`` in the output of the xref:{relative_path}snapmirror-
show.html[snapmirror show] command.

For relationships with ``_"Relationship Capability"_`` of ``_"8.2 and
above"_`` , you can track the progress of the operation using the
xref:{relative_path}snapmirror-show.html[snapmirror show] command.

For relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_``
, a job will be spawned to operate on the SnapMirror relationship, and the
job id will be shown in the command output. The progress of the job can be
tracked using the xref:{relative_path}job-show.html[job show] and
xref:{relative_path}job-history-show.html[job history show] commands.

For Vserver SnapMirror relationships, the `snapmirror update` command
```

```
makes the destination Vserver an up-to-date mirror of the source Vserver.

The `snapmirror update` command must be used from the destination Vserver
or cluster.

== Parameters

{ ``[-S, -source-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Source Path::
This parameter specifies the source endpoint of the SnapMirror
relationship in one of four path formats. The normal format includes the
names of the Vserver (vserver) and/or volume (volume). To support
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , a
format which also includes the name of the cluster (cluster) is provided.
The ``_"Pre 8.2"_`` format cannot be used when operating in a Vserver
context on relationships with ``_"Relationship Capability"_`` of ``_"8.2
and above"_`` . For SnapMirror relationships with an AltaVault source, the
source endpoint is specified in the form ``_hostip:/share/share-name_`` .
For SnapMirror relationships with a SolidFire source, the source endpoint
is specified in the form ``_hostip:/lun/name_`` .

| ``[-source-cluster <Cluster name>]`` - Source Cluster::
Specifies the source cluster of the SnapMirror relationship. If this
parameter is specified, the `-source-vserver` and `-source-volume`
parameters must also be specified. This parameter is only applicable for
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .
This parameter cannot be specified when operating in a Vserver context on
relationships with ``_"Relationship Capability"_`` of ``_"8.2 and
above"_`` .

``[-source-vserver <vserver name>]`` - Source Vserver::
Specifies the source Vserver of the SnapMirror relationship. For
relationships with volumes as endpoints, if this parameter is specified,
parameters `-source-volume` and for relationships with ``_"Relationship
Capability"_`` of ``_"Pre 8.2"_`` , `-source-cluster` must also be
specified. This parameter is not supported for relationships with non-
ONTAP source endpoints.

``[-source-volume <volume name>]`` - Source Volume }::
Specifies the source volume of the SnapMirror relationship. If this
parameter is specified, parameters `-source-vserver` and for relationships
with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-source-
cluster` must also be specified. This parameter is not supported for
relationships with non-ONTAP source endpoints.
```

```
{ ``-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}`` - Destination
Path::
This parameter specifies the destination endpoint of the SnapMirror
relationship in one of four path formats. The normal format includes the
names of the Vserver (vserver) and/or volume (volume). To support
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , a
format which also includes the name of the cluster (cluster) is provided.
The ``_"Pre 8.2"_`` format cannot be used when operating in a Vserver
context on relationships with ``_"Relationship Capability"_`` of ``_"8.2
and above"_`` . For SnapMirror relationships with an AltaVault
destinations, the destination endpoint is specified in the form
``_hostip:/share/share-name_`` . For SnapMirror relationships with a
SolidFire destinations, the destination endpoint is specified in the form
``_hostip:/lun/name_`` .

| ``[-destination-cluster <Cluster name>]`` - Destination Cluster::
Specifies the destination cluster of the SnapMirror relationship. If this
parameter is specified, parameters `-destination-vserver` and `-
destination-volume` must also be specified. This parameter is only
applicable for relationships with ``_"Relationship Capability"_`` of
``_"Pre 8.2"_`` . This parameter cannot be specified when operating in a
Vserver context on relationships with ``_"Relationship Capability"_`` of
``_"8.2 and above"_`` .

``-destination-vserver <vserver name>`` - Destination Vserver::
Specifies the destination Vserver of the SnapMirror relationship. For
relationships with volumes as endpoints, if this parameter is specified,
parameters `-destination-volume` and for relationships with
``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-destination-
cluster` must also be specified. This parameter is not supported for
relationships with non-ONTAP destination endpoints.

``-destination-volume <volume name>`` - Destination Volume }::
Specifies the destination volume of the SnapMirror relationship. If this
parameter is specified, parameters `-destination-vserver` and for
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` , `-
destination-cluster` must also be specified. This parameter is not
supported for relationships with non-ONTAP destination endpoints.

``[-s, -source-snapshot <text>]`` - Source Snapshot::
This optional parameter specifies a snapshot to transfer. The default
behavior for a data protection relationship with a read-write source is
that ONTAP creates a new snapshot and uses it as the basis for determining
what data is replicated; with this option, the specified snapshot will be
```

used instead. The default behavior for an extended data protection relationship depends on the relationship's policy type. For a data protection relationship, the specified snapshot must be newer than the latest common snapshot. For an extended data protection relationship, the specified snapshot can be newer or older than the common snapshot. This parameter is not supported for relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .

``[-k, -throttle <throttleType>]`` - Throttle (KB/sec)::
This optional parameter limits the network bandwidth used for the update transfer. It sets the maximum rate (in Kbytes/sec) at which data can be transferred during the operation. If this parameter is not specified, the throttle value configured for the relationship with the xref:{relative_path}snapmirror-create.html[snapmirror create] or xref:{relative_path}snapmirror-modify.html[snapmirror modify] command will be used. To fully use the network bandwidth available, set the throttle value to ``_unlimited_`` or ``_0_`` . The minimum throttle value is four Kbytes/sec, so if you specify a throttle value between ``_1_`` and ``_4_`` , it will be treated as if you specified ``_4_`` . For FlexGroup and Application Consistency Group relationships, the throttle value is applied individually to each constituent relationship. The `-throttle` parameter does not affect load-sharing transfers and transfers for other relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` confined to a single cluster.

``[-transfer-priority {low|normal}]`` - Transfer Priority::
This optional parameter specifies the priority at which the transfer runs. The default value for this parameter is the value in the SnapMirror policy associated with this relationship. This parameter is not applicable to relationships with a ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .

``[-enable-storage-efficiency <true>]`` - Enable Storage Efficient Transfers::
This is an optional parameter. For an extended data protection (XDP) relationship that is currently not storage efficient, set this parameter to true to enable storage efficient transfers. Storage efficient in this context refers to both over the wire efficiency and how the data is written to the destination volume. The transfer fails if storage efficiency cannot be achieved. If the transfer succeeds, future transfers will continue being storage efficient as long as it is still feasible, but will not fail if the transfer is not storage efficient. The default value is false. This parameter is not supported for relationships with non-Data ONTAP endpoints.

``[-w, -foreground <true>]`` - Foreground Process::
This specifies whether the operation runs as a foreground process. If this

```
parameter is specified, the default setting is ``_true_`` (the operation
runs in the foreground). When set to ``_true_`` , the command will not
return until the process completes. This parameter is only applicable to
relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_`` .

== Examples

To update the mirror relationship between the destination endpoint
``_vs2.example.com:dept_eng_dp_mirror3_`` and its source endpoint, type
the following command:
```

vs2.example.com::> snapmirror update -destination-path vs2.example.com:dept_eng_dp_mirror3

```
For relationships with ``_"Relationship Capability"_`` of ``_"Pre 8.2"_``
, to update the mirror relationship between the destination endpoint
``_cluster2://vs2.example.com/dept_eng_dp_mirror3_`` and its source
endpoint, type the following command:
```

cluster2::> snapmirror update -destination-path cluster2://vs2.example.com/dept_eng_dp_mirror3

```
To update the Vserver SnapMirror relationship between destination endpoint
``_dvs1.example.com:_`` and its source endpoint, type the following
command:
```

cluster2::> snapmirror update -destination-path dvs1.example.com:

```
To update the SnapMirror active sync relationship with the destination
Consistency Group ``_cg_dst_`` in Vserver ``_vs2.example.com_`` , type the
following command on the destination cluster:
```

destination::> snapmirror update -destination-path vs2.example.com:/cg/cg_dst

```
To update the Application Consistency Group relationship with the
destination Application Consistency Group ``_app_cg_dst_`` in Vserver
``_vs2.example.com_`` , type the following command on the destination
cluster:
```

destination::> snapmirror update -destination-path vs2.example.com:/cg/app_cg_dst

```
== Related Links

* xref:{relative_path}snapmirror-create.html[snapmirror create]
* xref:{relative_path}snapmirror-initialize.html[snapmirror initialize]
* xref:{relative_path}snapmirror-initialize-ls-set.html[snapmirror
initialize-ls-set]
* xref:{relative_path}snapmirror-update-ls-set.html[snapmirror update-ls-
set]
* xref:{relative_path}snapmirror-abort.html[snapmirror abort]
* xref:{relative_path}snapmirror-show.html[snapmirror show]
* xref:{relative_path}job-show.html[job show]
* xref:{relative_path}job-history-show.html[job history show]
* xref:{relative_path}snapmirror-modify.html[snapmirror modify]




= snapmirror config-replication commands

:leveloffset: +1


[[ID49b30a68efefe636eb292701844b4b0e]]
= snapmirror config-replication cluster-storage-configuration modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify SnapMirror storage configuration information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `snapmirror config-replication cluster-storage-configuration modify`
command modifies the configuration of storage used for configuration
replication.

== Parameters

``[-disallowed-aggregates <aggregate name>,...]`` - Disallowed Aggregates
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to set the list of storage aggregates that are not
```

```
available to host storage for configuration replication.

== Examples

The following example disallows two aggregates named ``_aggr1_`` and
``_aggr2_`` :
```

cluster1::*> snapmirror config-replication cluster-storage-configuration modify -disallowed-aggregates aggr1,aggr2

```
[[ID2d7ce5877144e5c65ec113ca9f517149]]
= snapmirror config-replication cluster-storage-configuration show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display SnapMirror storage configuration information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `snapmirror config-replication cluster-storage-configuration show`
command shows details of the configuration of the storage used for
configuration replication.

The information displayed is the following:

* Disallowed Aggregates - The list of storage aggregates that are
configured as not allowed to host storage areas.
* Auto-Repair - Displays ``_true_`` if the automatic repair of storage
areas used by configuration replication is enabled.
* Auto-Recreate - Displays ``_true_`` if the automatic recreation of
storage volumes used by configuration replication is enabled.
* Use Mirrored Aggregate - Displays ``_true_`` if storage areas for
configuration replication are to be hosted on a mirrored aggregate.



== Examples

The following is an example of the `snapmirror config-replication cluster-
storage-configuration show` command:
```

cluster1::*> snapmirror config-replication cluster-storage-configuration show
Disallowed Aggregates: -
Auto-Repair: true
Auto-Recreate: true
Use Mirrored Aggregate: true

```
[[ID5494714f3d952beb13a3f22322625e87]]
= snapmirror config-replication status show-aggregate-eligibility
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the SnapMirror configuration replication aggregate eligibility

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror config-replication status show-aggregate-eligibility`
command displays the SnapMirror configuration replication aggregate
eligibility.

== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
Display only rows that have a matching aggregate name.

``[-hosted-configuration-replication-volumes <volume name>,...]`` -
Currently Hosted Configuration Replication Volumes::
Display only rows that have matching configuration replication volumes
hosted on this aggregate.

``[-is-eligible-to-host-additional-volumes {true|false}]`` - Eligibility
to Host Another Configuration Replication Volume::
Display only rows that have a matching eligibility of the aggregate to
host additional configuration replication volumes.

``[-comment <text>]`` - Comment for Eligibility Status::
Display only rows that have a matching comment regarding the eligibility
of the aggregate to host configuration replication volumes.

== Examples
```

```
The following example shows the execution of the command in a SnapMirror
configuration with thirteen aggregates in the cluster:
```

clusA::snapmirror config-replication status> show-aggregate-eligibility
Eligible to
Aggregate Hosted Config Replication Vols Host Addl Vols Comments
---------- ---------- ---------- --------
a0 - false Root Aggregate
a1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_A true -
a2 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_B true -
a3 - false Unable to determine available space of aggregate
a4 - false Non-Local Aggregate
a5 - false Non-Home Aggregate
a6 - false Unable to determine mirror configuration
a7 - false Mirror configuration does not match requirement
a8 - false Disallowed Aggregate
a9 - false Insufficient Space - 10GB required
a10 - false Aggregate Offline
a11 - false Inconsistent Aggregate
a12 - false Aggregate Full
13 entries were displayed.

```
[[ID4127aee8e8a61d0dd42067c3321a11b4]]
= snapmirror config-replication status show-communication
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display SnapMirror configuration replication communication status
information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror config-replication status show-communication` command
displays the current SnapMirror configuration replication communication
status.

The command displays the following aspects of SnapMirror configuration
replication for each peer cluster:

* Remote Heartbeat: Verifies that the SnapMirror configuration replication
heartbeat with the remote cluster is healthy.
* Last Heartbeat Sent: Prints the timestamp of the last SnapMirror
```

```
configuration replication heartbeat sent to the remote cluster.
* Last Heartbeat Received: Prints the timestamp of the last SnapMirror
configuration replication heartbeat received from the remote cluster.



Additional information about the warnings (if any) and recovery steps can
be viewed by running the command with the -instance option.

== 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.

``[-cluster-uuid <UUID>]`` - Remote Cluster::
Display only rows that have a matching peer cluster UUID.

``[-cluster <text>]`` - Peer Cluster Name::
Display only rows that have matching peer cluster name.

``[-remote-heartbeat {ok|warning|not-run|not-applicable}]`` - Remote
Heartbeat::
Display only rows that have a matching remote heartbeat status.

``[-last-heartbeat-sent <MM/DD/YYYY HH:MM:SS>]`` - Last Heartbeat Sent
Time::
Display only rows that have a matching timestamp of the last heartbeat
sent.

``[-last-heartbeat-received <MM/DD/YYYY HH:MM:SS>]`` - Last Heartbeat
Received Time::
Display only rows that have a matching timestamp of the last heartbeat
received.

``[-heartbeat-recovery-steps <text>]`` - Heartbeat Recovery Steps::
Display only rows that have matching heartbeat recovery steps.

== Examples

The following example shows the execution of the command in a SnapMirror
configuration with two peer clusters:
```

```
clus1::*> snapmirror config-replication status show-communication
Peer Cluster: clus2
Remote Heartbeat: ok
Last Heartbeat Sent: 11/11/2014 11:11:45
Last Heartbeat Received: 11/11/2014 11:11:46
Peer Cluster: clus3
Remote Heartbeat: ok
Last Heartbeat Sent: 11/11/2014 11:11:26
Last Heartbeat Received: 11/11/2014 11:11:27

2 entries were displayed.
```

```
[[IDa24b7bac984eed467e58c1c255d18ae3]]
= snapmirror config-replication status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display SnapMirror configuration replication status information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror config-replication status show` command displays the
current SnapMirror configuration replication status.

The command displays the following aspects of SnapMirror configuration
replication:

* Enabled: Verifies that SnapMirror configuration replication is enabled
on the cluster.
* Running: Verifies that SnapMirror configuration replication is running
on the cluster.
* Storage Status: Verifies that SnapMirror configuration replication
storage is healthy.
* Storage In Use: Prints the location of SnapMirror configuration
replication storage.
* Storage Remarks: Prints the underlying root cause for non-healthy
SnapMirror configuration storage.
* Vserver Streams: Verifies that SnapMirror configuration replication
Vserver streams are healthy.
```

```
Additional information about the warnings (if any) and recovery steps can
be viewed by running the command with the -instance option.

== Parameters

``[-instance  ]``::
 If you specify the -instance parameter, the command displays detailed
information about all fields.

== Examples

The following example shows the execution of the command:
```

clus1::*> snapmirror config-replication status show
Enabled: true
Running: true
Storage Status: ok
Storage In Use: Cluster-wide Volume: MDV_CRS_3d47e9106b7d11e4a77b000c29f810a2_A
Storage Remarks: -
Vserver Streams: ok

```
:leveloffset: -1


= snapmirror failover commands

:leveloffset: +1


[[ID06a9f20c78c5d30ffa62e2c8e60f779a]]
= snapmirror failover show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display failover status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror failover show` command is used to monitor the progress of
the failover operation on the SnapMirror active sync relationship. You
must execute this command from the destination cluster of the SnapMirror
```

```
active sync relationship.

== 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.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Start Time::
Start operation time of the failover.

``[-source-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Source Path::
Specifies the source endpoint of the SnapMirror relationship.

``[-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Destination
Path::
Specifies the destination endpoint of the SnapMirror relationship.

``[-status {started|failed|completed|completed_with_warnings}]`` -
Failover Status::
Ths status of SnapMirror failover.

``[-error-reason <text>]`` - Error Reason::
The reason for failure occurred during SnapMirror failover.

``[-end-time <MM/DD/YYYY HH:MM:SS>]`` - End Time::
End operation time of the failover.

``[-primary-data-cluster <text>]`` - Primary Data Cluster::
Primary data cluster.

``[-last-progress-update <text>]`` - Last Progress Update::
Last progress update.

``[-type {planned|unplanned|incapable}]`` - Failover Type::
Whether the SnapMirror failover was planned or unplanned.

``[-error-reason-codes <integer>,...]`` - Error Reason codes::
```

```
The reason codes for failure occurred during SnapMirror failover.

== Examples

To see the progress of the failover operation performed on an SnapMirror
active sync relationship with destination Consistency Group ``_cg_dst_``
in Vserver ``_vs2.example.com_`` , type the following command from the
destination cluster:
```

destination::> snapmirror failover show -destination-path vs2.example.com:/cg/cg_dst

```
[[ID45be4522e0fb2a4ef62ed3ede173fbce]]
= snapmirror failover start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start planned failover for SnapMirror relationships with 'automated-
failover' policy-type

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror failover start` command initiates a planned failover
operation to change the roles of copies in a SnapMirror active sync
relationship without disrupting the client operations. The primary
writable volumes become the secondary data protection volumes while the
secondary data protection volumes become the primary writable volumes. The
new primary volumes can begin processing I/O requests locally without
disrupting the client operations. Asymmetric Logical Unit Access (ALUA)
reporting will also change as a part of this role change.

You should execute the `snapmirror failover start` command from the
destination cluster.

Make sure the relationship status is InSync and the ONTAP Mediator is
configured, connected, and in quorum before using this command to switch
the primary and secondary roles.

== Parameters
```

```
``-destination-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}`` - Destination
Path::
Specifies the destination endpoint of the SnapMirror relationship.

== Examples

To perform a planned failover operation of an SnapMirror active sync
relationship with source Consistency Group ``_cg_src_`` in Vserver
``_vs1.example.com_`` , and the destination Consistency Group ``_cg_dst_``
in Vserver ``_vs2.example.com_`` , type the following command from the
destination cluster:
```

destination::> snapmirror failover start -destination-path vs2.example.com:/cg/cg_dst

```
Note: The SnapMirror active sync relationship information in the
xref:{relative_path}snapmirror-show.html[snapmirror show] command will now
be available on the new destination that is Vserver ``_vs2.example.com_``
with ``_cg_dst_`` as the new source and ``_cg_src_`` as the new
destination Consistency Groups after the SnapMirror failover operation is
completed.

== Related Links

* xref:{relative_path}snapmirror-show.html[snapmirror show]




:leveloffset: -1


= snapmirror mediator commands

:leveloffset: +1


[[ID5941f7402cf9a22072b7f9ada5961a5f]]
= snapmirror mediator add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Create mediator config entry

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror mediator add` command configures the connection between a
pair of clusters and an ONTAP Mediator. It is mandatory to initialize the
ONTAP Mediator on one of the cluster peers before the SnapMirror active
sync relation with policy type ``_automated-failover_`` performs a planned
or an unplanned failover. You can initialize the ONTAP Mediator from
either cluster. When you issue the ONTAP Mediator add command on one
cluster, the ONTAP Mediator is automatically added on the other cluster.

== Parameters

``-mediator-address <IP Address>`` - Mediator Ip Address::
Specifies the IP address of the mediator.

``-peer-cluster <text>`` - Peer cluster::
Peer cluster with AutomatedFailOver SnapMirror relationships.

``-username <text>`` - Username::
User account at the mediator.

``[-port-number <integer>]`` - Port Number::
This optional parameter specifies the mediator service port number to
communicate with the mediator. The port number must be in the range 1025
to 65535 inclusive. The default port number is 31784.

== Examples

 The following example configures the connection between a mediator and a
pair of clusters.
```

clusA::> snapmirror mediator add -mediator-address 10.234.102.227 -peer-cluster clusB -username mediatoradmin
Notice: Enter the mediator password.
Enter the password:
Enter the password again:
Info: [Job: 114] 'mediator add' job queued

```
[[ID0ac45a49059caddd20ca4e12a062669d]]
= snapmirror mediator remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove mediator config entry

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror mediator remove` command deletes a mediator configuration
entry.

== Parameters

``-mediator-address <IP Address>`` - Mediator Ip Address::
Specifies the IP address of the mediator.

``-peer-cluster <text>`` - Peer cluster::
Peer cluster with AutomatedFailOver SnapMirror relationships.

== Examples

 The following example removes a mediator configuration entry.
```

clusA::> snapmirror mediator remove -mediator-address 10.140.102.227 -peer-cluster clusB
Info: [Job 36] 'mediator remove' job queued

```
[[ID53748b0e15118fb98bd507f380904869]]
= snapmirror mediator show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show mediator information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `snapmirror mediator show` command shows the status of the ONTAP
Mediator configuration.

== 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.

``[-mediator-address <IP Address>]`` - Mediator IP Address::
Ip address of the mediator.

``[-peer-cluster <text>]`` - Peer Cluster Name::
Peer cluster name.

``[-connection-status
{connected|unreachable|suspended|unusable|misconfigured|removing|not-
configured|unknown|adding}]`` - Connection Status::
Connection status of the cluster with the mediator.

``[-quorum-status {unknown|true|false}]`` - Mediator Quorum Status::
Shows whether the SnapMirror Consistency Group relationships are
synchronized with mediator.

``[-health-fetch-timeout <integer>]`` - Peer Health Fetch Timeout::
Timeout value (in seconds) for peer-cluster health fetch via mediator.

``[-connection-timeout <integer>]`` - Connection Timeout::
Timeout value (in seconds) for mediator connection.

== Examples

 The following example shows the list of mediator configurations.
```

clusA::*> snapmirror mediator show
Mediator Address Peer Cluster Connection Status Quorum Status
--- --- --- ---
10.140.102.227 clusB unreachable true
clusA::*> snapmirror mediator show -instance

Mediator Uuid: 416fbdee-c982-11e9-9034-005056a7124c
Mediator IP Address: 10.140.102.227
Peer Cluster: clusB
Peer Cluster Uuid: 771d9b13-c973-11e9-928e-005056a7a882
Connection Status: unreachable
Quorum Status: true
Health Fetch Timeout: 5
Connection Timeout: 5

```
[[ID77b4fd24a8fe8d1cd28f234fb6ed6f25]]
= snapmirror mediator primary-bias show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show Primary Bias Status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror mediator primary-bias show` command shows the status of
primary-bias. This command is applicable when mediator status is shown as
unreachable/misconfigured from both primary and secondary cluster.

This will also show the state corresponding to each CG endpoint residing
on the given cluster. Slave CG endpoints will show the state of "Failover
disabled". Master CG endpoints will show the state of Primary-bias
activation.

The status of primary-bias can be one of the following.  + primary-bias-
activated - represents the state that allows master endpoints to assume
the I/O serving authority + primary-bias-not-activated - represents the
state that mediator is reachable from primary and secondary clusters and
IO serving authority set to default + mediator-engaged-for-failover -
represents the state that mediator is engaged and Failover on slave
endpoints is possible + mediator-disengaged-for-failover - represents the
setting of the intermediate state preventing Failover on slave endpoints
of a CG +

== 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.

``[-cg-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Cg Path::
CG path.

``[-cg-role {destination|source}]`` - Role::
Indicates if CG is primary or secondary on local cluster.

``[-mb-state {Primary bias activated. Mediator Disengaged|Primary bias
not-activated. Mediator engaged|Mediator disengaged for failover|Mediator
engaged for failover}]`` - MB state::
Primary bias State.

== Examples

 The following example shows the status of primary bias
```

cluster1::> snapmirror mediator primary-bias show
CG Path Role Status
---------- ---------- ----------
vs1:/cg/dcg destination mediator-engaged-for-failover

```
[[ID9ba8d0f252a5aaf87cd369f97a244387]]
= snapmirror mediator primary-bias history show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show Primary Bias History

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror mediator primary-bias history show` command shows the
history of Primary bias events.
```

```
The setting/clearing of the Primary Bias states (including the
intermediate state) will be logged as event history per CG endpoint in the
"snapmirror mediator primary bias history" command. The "mediator-
engaged/disengaged-for-failover" actions will represent the
setting/clearing of the intermediate state preventing Failover on slave
endpoints of a CG. The "primary-bias-activated/not-activated" will
represent the state that allows master endpoints to assume the I/O serving
authority.

== 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.

``[-cg-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Cg Path::
CG path.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - TBA Task Start Time::
Time of event

``[-action {Primary bias activated. Mediator Disengaged|Primary bias not-
activated. Mediator engaged|Mediator disengaged for failover|Mediator
engaged for failover}]`` - Event::
Primary bias event.

== Examples

 The following example shows the history output
```

```
C1_sti7-vsim-ucs569m_cluster::*> snapmirror mediator primary-bias history show -cg-path vs0:/cg/scg
CG Path Time Action
------------- --------------- ------------
vs0:/cg/scg 7/25/2022 10:33:28
primary-bias-activated
vs0:/cg/scg 7/25/2022 22:42:33
primary-bias-not-activated
2 entries were displayed.
C1_sti7-vsim-ucs569m_cluster::*>
```

```
[[ID9d41614cf60fe2a6f6f0218e6fbfd22b]]
= snapmirror mediator tba-history show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show Mediator Agent Event History

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `snapmirror mediator tba-history show` command shows the history of
the ONTAP Mediator Agent events..

== 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.

``[-cg-rel-uuid <UUID>]`` - Cg Relationship Uuid::
Cg Relationship Uuid.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - TBA Task Start Time::
Mediator Agent Task Start Time.

``[-token-num <integer>]`` - Token Number::
Token Number.

``[-cg-path
{<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<h
ostip:/share/share-name>|<[vserver:]/cg/[app-cgname]>}]`` - Cg Path::
Cg Path.

``[-aux-generation-num <integer>]`` - Auxillary Generation Number::
Auxiliary Generation Number.
```

```
``[-cg-rel-task <text>]`` - TBA Task Type::
Mediator Agent Task Type.

``[-cg-rel-state {invalid-state|waiting|task-processing|task-complete}]``
- Cg Management State::
Cg Management State.

``[-end-time <MM/DD/YYYY HH:MM:SS>]`` - TBA Task End Time::
Mediator Agent Task End Time.

``[-cg-rel-action {invalid|none|continue|stop|failover-incapable|auto-
failover|success|failed|mismatch|retry|success-pending|stop-
pending|continue-pending|failover-incapable-pending|try-failover-
incapable-ic|try-failover-incapable-med|try-auto-failover}]`` - Cg
Management Action::
Cg Management Action.

``[-errCode <integer>]`` - Error Code for Task Failure::
Error Code for Task Failure.

== Examples

 The following example shows the list of mediator configurations.
```

C1_sti89-vsim-ucs525q_cluster::> snapmirror mediator tba-history show
(snapmirror mediator tba-history show)
Path Token API Event Start Time State End Time Action
---------- ----- ----- ------- ---------- ---------- ---------- ----------
vs0:/cg/cg1_src 5 4 Process Set Resync Context 4/28/2021 11:09:52 task-complete 4/28/2021 11:09:52 continue
vs0:/cg/cg1_src 6 1 Process Master InSync 4/28/2021 11:10:08 task-complete 4/28/2021 11:10:08 success
2 entries were displayed.
C1_sti89-vsim-ucs525q_cluster::> snapmirror mediator tba-history show -instance
Cg Relationship Uuid: ad49a4f1-a833-11eb-9846-005056a7c906
Mediator Agent Task Start Time: 4/28/2021 11:09:52
Token Number: 5
Cg Path: vs0:/cg/cg1_src
Auxillary Generation Number: 4
Mediator Agent Task Type: Process Set Resync Context
Cg Management State: task-complete
Mediator Agent Task End Time: 4/28/2021 11:09:52
Cg Management Action: continue
Error Code for Task Failure: 0
Cg Relationship Uuid: ad49a4f1-a833-11eb-9846-005056a7c906
Mediator Agent Task Start Time: 4/28/2021 11:10:08
Token Number: 6
Cg Path: vs0:/cg/cg1_src
Auxillary Generation Number: 1
Mediator Agent Task Type: Process Master InSync
Cg Management State: task-complete

Mediator Agent Task End Time: 4/28/2021 11:10:08
Cg Management Action: success
Error Code for Task Failure: 0

```
:leveloffset: -1


= snapmirror object-store commands

:leveloffset: +1


[[IDcdc715c58153dd7385278d3eddcf6f2d]]
= snapmirror object-store config create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Define the configuration for a SnapMirror object store

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror object-store config create` command is used by a cluster
administrator to tell ONTAP how to connect to an object store. Following
pre-requisites must be met before creating an object store configuration
in ONTAP.

* A valid data bucket or container must be created with the object store
provider. This assumes that the user has valid account credentials with
the object store provider to access the data bucket.
* The ONTAP node must be able to connect to the object store. This
includes

  * Fast, reliable connectivity to the object store.
  * An inter-cluster LIF (Logical Interface) must be configured on the
cluster.
  * If SSL/TLS authentication is required, then valid certificates must be
installed.
```

```
An object-store configuration once created must not be reassociated with a
different object-store or container. See xref:{relative_path}snapmirror-
object-store-config-modify.html[snapmirror object-store config modify]
command for more information.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the vserver on which the object store
configuration needs to be created.

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter specifies the name that will be used to identify the object
store configuration. The name can contain the following characters: "_",
"-", A-Z, a-z, and 0-9. The first character must be one of the following:
"_", A-Z, or a-z.

``-usage {data|metadata}`` - Object Store Use::
This parameter specifies the usage for an object store configuration.

``-provider-type <providerType>`` - Type of the Object Store Provider::
This parameter specifies the type of object store provider that will be
attached to the aggregate. Valid options are: AWS_S3 (Amazon S3 storage),
Azure_Cloud (Microsoft Azure Cloud), SGWS (StorageGrid WebScale), IBM_COS
(IBM Cloud Object Storage), AliCloud (Alibaba Cloud Object Storage
Service), GoogleCloud (Google Cloud Storage) and ONTAP_S3.

``-server <Remote InetAddress>`` - Fully Qualified Domain Name of the
Object Store Server::
This parameter specifies the Fully Qualified Domain Name (FQDN) of the
remote object store server. For Amazon S3, server name must be an AWS
regional endpoint in the format s3.amazonaws.com or s3-
<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The
region of the server and the bucket must match. For more information on
AWS regions, refer to 'Amazon documentation on AWS regions and endpoints'.
For Azure, if the `-server` is a "blob.core.windows.net" or a
"blob.core.usgovcloudapi.net", then a value of `-azure-account` followed
by a period will be added in front of the server.

``[-is-ssl-enabled {true|false}]`` - Is SSL/TLS Enabled::
This parameter indicates whether a secured SSL/TLS connection will be used
during data access to the object store. The default value is `true` .

``[-port <integer>]`` - Port Number of the Object Store::
This parameter specifies the port number on the remote server that ONTAP
```

```
will use while establishing connection to the object store.

``-container-name <text>`` - Data Bucket/Container Name::
This parameter specifies the data bucket or container that will be used
for read and write operations.
+
NOTE: This name cannot be modified once a configuration is created.
+



``[-access-key <text>]`` - Access Key ID for S3 Compatible Provider
Types::
This parameter specifies the access key (access key ID) required to
authorize requests to the AWS S3, SGWS, IBM COS object stores and
ONTAP_S3. For an Azure object store see `-azure-account.`

``[-ipspace <IPspace>]`` - IPspace to Use in Order to Reach the Object
Store::
This optional parameter specifies the IPspace to use to connect to the
object store. Default value: ``_Default_`` .

``[-use-iam-role {true|false}]`` - (DEPRECATED)-Use IAM Role for AWS Cloud
Volumes ONTAP::
This optional parameter is deprecated. Use `-auth-type` instead. Note,
that `-auth-type EC2-IAM` is an equivalent of `-use-iam-role true` , and
`-auth-type key` is an equivalent of `-use-iam-role false` .

``[-secret-password <text>]`` - Secret Access Key for S3 Compatible
Provider Types::
This parameter specifies the password (secret access key) to authenticate
requests to the AWS S3, SGWS, IBM COS object stores and ONTAP_S3. If the
`-access-key` is specified but the `-secret-password` is not, then one
will be asked to enter the `-secret-password` without echoing the input.

``[-is-certificate-validation-enabled {true|false}]`` - Is SSL/TLS
Certificate Validation Enabled::
This parameter indicates whether an SSL/TLS certificate of an object store
server is validated whenever an SSL/TLS connection to an object store
server is established. This parameter is only applicable when `is-ssl-
enabled` is `true` . The default value is `true` . It is recommended to
use the default value to make sure that ONTAP connects to a trusted object
store server, otherwise identities of an object store server are not
verified.

``[-azure-account <text>]`` - Azure Account::
This parameter specifies the account required to authorize requests to the
```

```
Azure object store. For other object store providers see access-key.
+
NOTE: The value of this field cannot be modified once a configuration is
created.
+


``[-ask-azure-private-key {true|false}]`` - Ask to Enter the Azure Access
Key without Echoing::
If this parameter is true then one will be asked to enter `-azure-private-
key` without echoing the input. Default value: ``_true_`` .

``[-azure-private-key <text>]`` - Azure Access Key::
This parameter specifies the access key required to authenticate requests
to the Azure object store. See also `ask-azure-private-key` . For other
object store providers see `-secret-password` .

``[-server-side-encryption {none | SSE-S3}]`` - Encryption of Data at Rest
by the Object Store Server::
This parameter specifies if AWS or other S3 compatible object store server
must encrypt data at rest. The available choices depend on provider-type.
`none` encryption (no encryption required) is supported by all types of S3
(non-Azure) object store servers. `SSE-S3` encryption is supported by and
is a default for all types of S3 (non-Azure) object store servers except
ONTAP_S3. This is an advanced property. In most cases it is best not to
change default value of "sse_s3" for object store servers which support
SSE-S3 encryption. The encryption is in addition to any encryption done by
ONTAP at a volume or at an aggregate level.

``[-url-style {path-style | virtual-hosted-style}]`` - URL Style Used to
Access S3 Bucket::
This parameter specifies the URL style used to access S3 bucket. This
option is only available for non-Azure object store providers. The
available choices and default value depend on provider-type.

``[-iamra-session-token <text>]`` - IAMRA Session Token for
Authentication::
This parameter specifies a temporary token for S3 SnapMirror which will
expire periodically. This will increase security.

``[-azure-msi-token <text>]`` - Azure MSI Token for Authentication::
This optional parameter specifies the Managed Service Identity (MSI) token
required to authenticate to the Azure object store. This form of
authentication is only supported on Azure NetApp Files.

== Examples
```

```
The following example creates a cluster scoped object store configuration
named ``_objectStoreName._``
```

cluster1::*> snapmirror object-store config create
-object-store-name objectStoreName -usage data -owner snapmirror
-provider-type SGWS -server objectStoreServer.com
-container-name containerName -is-ssl-enabled true
-is-certificate-validation-enabled false
-ipspace Default -access-key userAccessKey
-secret-password userSecretPassWord

```
== Related Links

* xref:{relative_path}snapmirror-object-store-config-
modify.html[snapmirror object-store config modify]




[[ID9b0d25d1816d0fb9dc02fa457f725d17]]
= snapmirror object-store config delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete SnapMirror object store configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror object-store config delete` command removes an existing
object store configuration in ONTAP.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the vserver name on which the object store
configuration has been configured.

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter specifies the object store configuration to be deleted.

== Examples

The following example deletes an object store configuration named
``_objectStoreName._``
```

cluster1::*> snapmirror object-store config delete
-object-store-name objectStoreName

```
[[IDc78272db27e6c2ae4570e7a3f14475b7]]
= snapmirror object-store config modify
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify SnapMirror object store configuration attributes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The xref:{relative_path}storage-aggregate-object-store-config-
modify.html[storage aggregate object-store config modify] command is used
to update one or more of object store configuration parameters.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the vserver on which the object store
configuration needs to be created.

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter identifies the configuration to be modified.

``[-usage {data|metadata}]`` - Object Store Use::
This parameter specifies the usage for an object store configuration.

``[-new-object-store-name <text>]`` - Object Store Configuration New
Name::
This optional parameter specifies the new name for the object store
configuration.

``[-server <Remote InetAddress>]`` - Fully Qualified Domain Name of the
Object Store Server::
This optional parameter specifies the new Fully Qualified Domain Name
(FQDN) of the same object store server. For Amazon S3, server name must be
an AWS regional endpoint in the format s3.amazonaws.com or s3-
<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The
region of the server and the bucket must match. For more information on
AWS regions, refer to 'Amazon documentation on AWS regions and endpoints'.
For Azure, if the `-server` is a "blob.core.windows.net" or a
"blob.core.usgovcloudapi.net", then the value of azure-account in the
configuration followed by a period will be added in front of the server.
Note that the value of `azure-account` cannot be modified.
```

```
``[-is-ssl-enabled {true|false}]`` - Is SSL/TLS Enabled::
This optional parameter indicates whether a secured SSL/TLS connection
will be used during data access to the object store.

``[-port <integer>]`` - Port Number of the Object Store::
This optional parameter specifies a new port number to connect to the
object store server indicated in the `-server` parameter.

``[-access-key <text>]`` - Access Key ID for S3 Compatible Provider
Types::
This optional parameter specifies a new access key (access key ID) for the
AWS S3, SGWS, IBM COS object stores and ONTAP S3.

``[-ipspace <IPspace>]`` - IPspace to Use in Order to Reach the Object
Store::
This optional parameter specifies new ipspace values for the
configuration.

``[-use-iam-role {true|false}]`` - (DEPRECATED)-Use IAM Role for AWS Cloud
Volumes ONTAP::
This optional parameter is deprecated. Use `-auth-type` instead. Note,
that `-auth-type EC2-IAM` is an equivalent of `-use-iam-role true` , and
`-auth-type key` is an equivalent of `-use-iam-role false` .

``[-secret-password <text>]`` - Secret Access Key for S3 Compatible
Provider Types::
This optional parameter specifies a new password (secret access key) for
the AWS S3, SGWS, IBM COS object stores and ONTAP S3. For an Azure object
store see `-azure-private-key` . If the `-access-key` is specified but the
`-secret-password` is not then one will be asked to enter the `-secret-
password` without echoing the input.

``[-is-certificate-validation-enabled {true|false}]`` - Is SSL/TLS
Certificate Validation Enabled::
This optional parameter indicates whether an SSL/TLS certificate of an
object store server is validated whenever an SSL/TLS connection to an
object store server is established. This parameter is only applicable when
`is-ssl-enabled` is `true` . It is recommended to keep the default value
which is `true` to make sure that ONTAP connects to a trusted object store
server, otherwise identities of an object store server are not verified.

``[-ask-azure-private-key {true|false}]`` - Ask to Enter the Azure Access
Key without Echoing::
If this optional parameter is true then one will be asked to enter the `-
azure-private-key` without echoing the input.
```

```
``[-azure-private-key <text>]`` - Azure Access Key::
This optional parameter specifies a new access key for Azure object store.
For other object store providers see `secret-password` . See also `ask-
azure-private-key` .

``[-server-side-encryption {none | SSE-S3}]`` - Encryption of Data at Rest
by the Object Store Server::
This parameter specifies if AWS or other S3 compatible object store server
must encrypt data at rest. The available choices depend on provider-type.
`none` encryption (no encryption required) is supported by all S3 (non-
Azure) object store servers. `SSE-S3` encryption is supported by all S3
(non-Azure) object store servers except ONTAP_S3. This is an advanced
property. In most cases it is best not to change default value of "sse_s3"
for object store servers which support SSE-S3 encryption. The encryption
is in addition to any encryption done by ONTAP at a volume or at an
aggregate level. Note that changing this option does not change encryption
of data which already exist in the object store.

``[-url-style {path-style | virtual-hosted-style}]`` - URL Style Used to
Access S3 Bucket::
This parameter specifies the URL style used to access S3 bucket. This
option is only available for non-Azure object store providers. The
available choices and default value depend on provider-type.

``[-iamra-session-token <text>]`` - IAMRA Session Token for
Authentication::
This parameter specifies a temporary token for S3 SnapMirror which will
expire periodically. This will increase security.

``[-azure-msi-token <text>]`` - Azure MSI Token for Authentication::
This optional parameter specifies the new Managed Service Identity (MSI)
token required to authenticate to the Azure object store. This form of
authentication is only supported on Azure NetApp Files.

== Examples

The following example modifies an object-store configuration named
``_objectStoreName_`` to a new name ``_newName._``
```

cluster::*> snapmirror object-store config modify
-object-store-name objectStoreName
-new-object-store-name newName

```
== Related Links

* xref:{relative_path}storage-aggregate-object-store-config-
```

```
modify.html[storage aggregate object-store config modify]




[[IDd04a6223adc0ea0fa9267727ef3e8294]]
= snapmirror object-store config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of SnapMirror object store configurations

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror object-store config show` command displays information
about all existing object store configurations in the cluster.

== 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.

``[-vserver <vserver name>]`` - Vserver Name::
If this parameter is specified, the command displays information only
about object store configurations which are configured within this
vserver.

``[-object-store-name <text>]`` - Object Store Configuration Name::
If this parameter is specified, the command displays information only
about object store configurations whose name matches the specified names.

``[-usage {data|metadata}]`` - Object Store Use::
Specifies what this objectc store configuration is used for.

``[-vsUuid <UUID>]`` - Vserver UUID::
```

```
Vserver UUId.

``[-config-id <integer>]`` - Object Store Configuration ID::
If this parameter is specified, the command displays information only
about object store configurations whose configuration ID matches the
specified value.

``[-provider-type <providerType>]`` - Type of the Object Store Provider::
If this parameter is specified, the command displays information only
about object store configurations whose provider type matches the
specified value.

``[-server <Remote InetAddress>]`` - Fully Qualified Domain Name of the
Object Store Server::
If this parameter is specified, the command displays information only
about object store configurations whose server name matches the specified
value. The server name is specified as a Fully Qualified Domain Name
(FQDN).

``[-is-ssl-enabled {true|false}]`` - Is SSL/TLS Enabled::
If this parameter is specified, the command displays information only
about object store configurations whose status about the use of secured
communication over the network matches the specified value.

``[-port <integer>]`` - Port Number of the Object Store::
If this parameter is specified, the command displays information only
about object store configurations whose port numbers matches the specified
value.

``[-container-name <text>]`` - Data Bucket/Container Name::
If this parameter is specified, the command displays information only
about object store configurations whose container name matches the
specified value. ONTAP uses this container name or object store data
bucket while accessing data from the object store.

``[-access-key <text>]`` - Access Key ID for S3 Compatible Provider
Types::
If this parameter is specified, the command displays information only
about AWS S3, SGWS, IBM COS object store configurations and ONTAP S3 whose
access key matches the specified value. ONTAP requires the access key for
authorized access to the object store.

``[-ipspace <IPspace>]`` - IPspace to Use in Order to Reach the Object
Store::
If this parameter is specified, the command displays information only
about object store configurations whose IPspace matches the specified
```

```
value. ONTAP uses the IPspace value to connect to the object store.

``[-use-iam-role {true|false}]`` - (DEPRECATED)-Use IAM Role for AWS Cloud
Volumes ONTAP::
If this parameter is specified, the command displays information only
about object store configurations whose IAM role status flag matches the
specified value. The `-iam-role` and `-use-iam-role` parameters are
relevant only in the context of AWS object store and indicates whether IAM
role must be used for accessing it. The IAM credentials can be obtained
only through AWS Cloud Volumes ONTAP.

``[-is-certificate-validation-enabled {true|false}]`` - Is SSL/TLS
Certificate Validation Enabled::
If this parameter is specified, the command displays information only
about object store configurations whose status about the validation of
SSL/TLS certificate matches the specified value.

``[-azure-account <text>]`` - Azure Account::
If this parameter is specified, the command displays information only
about Azure object store configurations whose account matches the
specified value. ONTAP requires the Azure account for authorized access to
the Azure object store.

``[-server-side-encryption {none | SSE-S3}]`` - Encryption of Data at Rest
by the Object Store Server::
If this parameter is specified, the command displays information only
about object store configurations whose server-side encryption matches the
specified value.

``[-url-style {path-style | virtual-hosted-style}]`` - URL Style Used to
Access S3 Bucket::
If this parameter is specified, the command displays information only
about object store configurations whose URL style matches the specified
value.

== Examples

The following example displays information about all existing object store
configurations in the cluster.
```

cluster1::*> snapmirror object-store config show

```
:leveloffset: -1


= snapmirror policy commands
```

```
:leveloffset: +1


[[IDbb94655d1ea70ec1014358dc3f55f772]]
= snapmirror policy add-rule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a new rule to SnapMirror policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror policy add-rule` command adds a rule to a SnapMirror
policy. Rules define which snapshots are protected by vault relationships
or define the schedule at which snapshots are created on the SnapMirror
destination. Rules which do not include a schedule are rules for
protecting snapshots. Rules which include a schedule are rules for
creating snapshots on the SnapMirror destination. A rule with a schedule
can only be added to SnapMirror policies of type ``_vault_`` or ``_mirror-
vault_`` . A rule must not be added to a policy that will be associated
with a SnapMirror data protection relationship. A policy that will be
associated with a SnapMirror vault relationship must have at least one
rule and at most ten rules. A SnapMirror policy with rules must have at
least one rule without a schedule.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver for the SnapMirror policy.

``-policy <sm_policy>`` - SnapMirror Policy Name::
Specifies the SnapMirror policy name.

``-snapmirror-label <text>`` - Snapshot Copy Label::
This parameter is primarily used for the purpose of snapshot selection for
extended data protection (XDP) relationships. Only snapshots that have a
SnapMirror label that matches this parameter will be transferred to the
SnapMirror destination. However, when this parameter is associated with a
rule containing a schedule, snapshots will be created on the SnapMirror
destination using this snapmirror-label parameter. The label can be 31 or
```

fewer characters in length. SnapMirror policies of type ``_async-mirror_`` and ``_mirror-vault_`` have a rule added for label ``_sm_created_`` at the time of policy creation. This rule cannot be removed or modified by the user. This rule when coupled with ``_create-snapshot_`` field set to ``_true_`` indicates that the SnapMirror relationship using this policy shall create a new snapshot and transfer it as part of a xref:{relative_path}snapmirror-update.html[snapmirror update] operation. SnapMirror policies of type ``_async-mirror_`` support one additional rule with SnapMirror label ``_all_source_snapshots_`` . This rule along with the rule for SnapMirror label ``_sm_created_`` indicates that ``_all_`` new snapshots on the primary volume along with the newly created snapshot are transferred as a part of a xref:{relative_path}snapmirror-update.html[snapmirror update] or xref:{relative_path}snapmirror-initialize.html[snapmirror initialize] operation. Rules with any other SnapMirror labels cannot be added to SnapMirror policies of type ``_async-mirror_`` . The rule for label ``_sm_created_`` when added to a `snapmirror policy` of type ``_vault_`` indicates that all SnapMirror created snapshots of the primary volume are selected for transfer.

``-keep <text>`` - Snapshot Copy Retention Count::
Specifies the maximum number of snapshots that are retained on the SnapMirror destination volume for a rule. The total number of snapshots retained for all the rules in a policy cannot exceed 1019. For all the rules in SnapMirror policies of type ``_async-mirror_`` , this parameter must be set to value ``_1_`` .

``[-preserve {true|false}]`` - Snapshot Copy Preserve Enabled::
Specifies the behavior when the snapshot retention count is reached on the SnapMirror vault destination for the rule. The default value is ``_false_`` , which means that the oldest snapshot will be deleted to make room for new ones only if the number of snapshots has exceeded the retention count specified in the "keep" parameter. When set to ``_true_`` , and when the snapshots have reached the retention count, then an incremental SnapMirror vault update transfer will fail or if the rule has a schedule, snapshots will no longer be created on the SnapMirror destination. For all the rules in SnapMirror policies of type ``_async-mirror_`` this parameter must be set to value ``_false_`` .

``[-warn <integer>]`` - Warning Threshold Count::
Specifies the warning threshold count for the rule. The default value is ``_0_`` . When set to a value greater than zero, an event is generated after the number of snapshots (for the particular rule) retained on a SnapMirror vault destination reaches the specified warn limit. The preserve parameter for the rule must be ``_true_`` to set the warn parameter to a value greater than zero.

```
``[-schedule <text>]`` - Snapshot Copy Creation Schedule::
This optional parameter specifies the name of the schedule associated with
a rule. This parameter is allowed only for rules associated with
SnapMirror policies of type ``_vault_`` or ``_mirror-vault_`` . When this
parameter is specified, snapshots are directly created on the SnapMirror
destination. The snapshots created will have the same content as the
latest snapshot already present on the SnapMirror destination. snapshots
on the source that have a SnapMirror label matching this rule will not be
selected for transfer. The default value is ``_-_`` .
+
NOTE: You define and name a schedule using the xref:{relative_path}job-
schedule-cron-create.html[job schedule cron create] command.
+


``[-prefix <text>]`` - Snapshot Copy Creation Prefix::
This optional parameter specifies the prefix for the snapshot name to be
created as per the schedule. If no value is specified, then the
``_snapmirror-label_`` will be used as the prefix. The prefix parameter
can only be specified for rules which have a schedule.

``[-retention-period {{<integer> seconds|minutes|hours|days|months|years}
| infinite}]`` - Snapshot Copy Retention Period::
This optional parameter specifies the duration for which the snapshots
associated with this rule will be locked. The parameter is specified as a
number followed by a suffix. The valid suffixes are ``_seconds_`` ,
``_minutes_`` , ``_hours_`` , ``_days_`` , ``_months_`` , and ``_years_``
. For example, a value of ``_6months_`` represents a retention period of 6
months. It can also be set to a special value of ``_infinite_`` for
infinite retention. The destination volume must have the ``_snapshot-
locking-enabled_`` parameter set to true for this field to take effect
else it will be ignored.

== Examples

The following example adds a rule named ``_nightly_`` to the SnapMirror
policy named ``_TieredBackup_`` on Vserver ``_vs0.example.com_`` . The
rule will retain a maximum of ``_5 nightly_`` snapshots.
```

vs0.example.com::> snapmirror policy add-rule -vserver vs0.example.com
-policy TieredBackup -snapmirror-label nightly -keep 5

```
== Related Links

* xref:{relative_path}snapmirror-update.html[snapmirror update]
* xref:{relative_path}snapmirror-initialize.html[snapmirror initialize]
```

```
* xref:{relative_path}job-schedule-cron-create.html[job schedule cron
create]




[[ID4d3e4bf661cfe92258e2e2c050bd483e]]
= snapmirror policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new SnapMirror policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror policy create` command creates a SnapMirror policy. When
applied to a SnapMirror relationship, the SnapMirror policy controls the
behavior of the relationship and specifies the configuration attributes
for that relationship. The policies ``_DPDefault_`` ,
``_MirrorAllSnapshots_`` , ``_MirrorAndVault_`` , ``_MirrorLatest_`` ,
``_Unified7year_`` , and ``_XDPDefault_`` are created by the system for
asynchronous replication. The policies ``_Sync_`` , ``_StrictSync_`` ,
``_AutomatedFailOver_`` , and ``_AutomatedFailOverDuplex_`` are created by
the system for synchronous replication.

NOTE: All SnapMirror policies have a field ``_create-snapshot_`` . This
field specifies whether SnapMirror creates a new snapshot on the primary
volume at the beginning of a xref:{relative_path}snapmirror-
update.html[snapmirror update] or xref:{relative_path}snapmirror-
resync.html[snapmirror resync] operation. Currently this field cannot be
set or modified by the user. It is set to ``_true_`` for SnapMirror
policies of type ``_async-mirror_`` and ``_mirror-vault_`` at the time of
creation. SnapMirror policies of type ``_vault_`` have ``_create-
snapshot_`` set to ``_false_`` at the time of creation.



NOTE: Use the xref:{relative_path}snapmirror-policy-add-
rule.html[snapmirror policy add-rule] command to add a rule to a policy.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver for the SnapMirror policy.

``-policy <sm_policy>`` - SnapMirror Policy Name::
This parameter specifies the SnapMirror policy name. A policy name can be
made up of the characters A to Z, a to z, 0 to 9, ".", "-", and "_". The
name can be up to 256 characters in length.

``[-type {vault|async-mirror|mirror-vault|strict-sync-mirror|sync-
mirror|automated-failover|automated-failover-duplex|continuous}]`` -
Snapmirror Policy Type::
This parameter specifies the SnapMirror policy type. The supported values
are ``_async-mirror_`` , ``_vault_`` , ``_mirror-vault_`` , ``_sync-
mirror_`` and ``_strict-sync-mirror_`` . Data protection (DP)
relationships support only ``_async-mirror_`` policy type, while extended
data protection (XDP) relationships support all policy types.
+
If the type is set to ``_async-mirror_`` then the policy is for Disaster
Recovery. When the policy type is associated with extended data protection
(XDP) relationships, xref:{relative_path}snapmirror-update.html[snapmirror
update] and xref:{relative_path}snapmirror-resync.html[snapmirror resync]
operations transfer selected snapshots from the primary volume to the
secondary volume. The selection of snapshots is governed by the rules in
the policy. However xref:{relative_path}snapmirror-
initialize.html[snapmirror initialize] and xref:{relative_path}snapmirror-
update.html[snapmirror update] operations on data protection (DP)
relationships ignore the rules in the policy and transfer ``_all_``
snapshots of the primary volume which are newer than the common snapshot
on the destination. For both data protection (DP) and extended data
protection (XDP) relationships, the snapshots are kept on the secondary
volume as long as they exist on the primary volume. Once a protected
snapshot is deleted from the primary volume, it is deleted from the
secondary volume as part of the next transfer. The policy type supports
rules with certain pre-defined label names only. Refer to the man page for
the xref:{relative_path}snapmirror-policy-add-rule.html[snapmirror policy
add-rule] command for the details.
+
If the type is set to ``_vault_`` then the policy is used for Backup and
Archive. The rules in this policy type determine which snapshots are
protected and how long they are retained on the secondary volume. This
policy type is supported by extended data protection (XDP) relationships
only.
+
```

```
If the type is set to ``_mirror-vault_`` then the policy is used for
unified data protection which provides both Disaster Recovery and Backup
using the same secondary volume. This policy type is supported by extended
data protection (XDP) relationships only.
+
If the type is set to ``_sync-mirror_`` or ``_strict-sync-mirror_`` then
the policy is used for synchronous Disaster Recovery. These are supported
only by extended data protection (XDP) relationships between FlexVol
volumes. Once the relationship is initialized with
xref:{relative_path}snapmirror-initialize.html[snapmirror initialize] ,
the relationship will be InSync such that all writes to the primary will
be replicated to the secondary before the write is acknowledged to the
client. Upon a replication failure, relationship falls OutOfSync. Upon an
OutOfSync event, the ``_strict-sync-mirror_`` variant restricts further
client IO on the primary, whereas the ``_sync-mirror_`` variant does not.
SnapMirror will automatically trigger resync to bring the OutOfSync
relationships back InSync as soon as it can, unless the relationship is
``_Quiesced_`` or ``_Broken-off_`` . Once a relationship is initialized,
you normally use the xref:{relative_path}snapmirror-
quiesce.html[snapmirror quiesce] command to stop synchronous replication
and the xref:{relative_path}snapmirror-resume.html[snapmirror resume]
command to resume synchronous replication. These policy types do not
support replication of user snapshots.

``[-comment <text>]`` - Comment::
Specifies a text comment for the SnapMirror policy. If the comment
contains spaces, it must be enclosed within quotes.

``[-tries <integer_or_unlimited>]`` - Tries Limit::
Determines the maximum number of times to attempt each manual or scheduled
transfer for a SnapMirror relationship. The value of this parameter must
be a positive integer or ``_unlimited_`` . The default value is ``_8_`` .

``[-transfer-priority {low|normal}]`` - Transfer Scheduling Priority::
Specifies the priority at which a transfer runs. The supported values are
``_normal_`` or ``_low_`` . The ``_normal_`` transfers are scheduled
before the ``_low_`` priority transfers. The default is ``_normal_`` .

``[-ignore-atime {true|false}]`` - Ignore File Access Time::
This parameter applies only to extended data protection (XDP)
relationships. It specifies whether incremental transfers will ignore
files which have only their access time changed. The supported values are
``_true_`` or ``_false_`` . The default is ``_false_`` .

``[-restart {always|never|default}]`` - Restart Behavior::
This parameter applies only to data protection relationships. It defines
```

```
the behavior of SnapMirror if an interrupted transfer exists. The
supported values are ``_always_`` , ``_never_`` , or ``_default_`` . If
the value is set to ``_always_`` , an interrupted SnapMirror transfer
always restarts provided it has a restart checkpoint and the conditions
are the same as they were before the transfer was interrupted. In
addition, a new SnapMirror snapshot is created which will then be
transferred. If the value is set to ``_never_`` , an interrupted
SnapMirror transfer will never restart, even if a restart checkpoint
exists. A new SnapMirror snapshot will still be created and transferred.
ONTAP version 8.2 will interpret a value of ``_default_`` as being the
same as ``_always_`` . Vault transfers will always resume based on a
restart checkpoint, provided the snapshot still exists on the source
volume.

``[-is-network-compression-enabled {true|false}]`` - Is Network
Compression Enabled::
Specifies whether network compression is enabled for transfers. The
supported values are ``_true_`` or ``_false_`` . The default is
``_false_`` .

``[-rpo <integer>]`` - Recovery Point Objective (seconds)::
Specifies the time for recovery point objective, in seconds. This
parameter is only supported for a policy of type ``_continuous_`` .

``[-always-replicate-snapshots {true|false}]`` - This prioritizes
replication of app-consistent snapshots over synchronous replication::
If this parameter is set to true, it specifies that SnapMirror Synchronous
relationships will lose the zero RPO protection upon failure in
replicating application created snapshots. The default value is false.

``[-common-snapshot-schedule <text>]`` - Common Snapshot Copy Creation
Schedule for SnapMirror Synchronous [.advaced-priv]#(privilege:
advanced)#::
Specifies the common snapshot creating schedule. This parameter is only
supported for SnapMirror Synchronous relationships.

``[-are-data-ops-sequentially-split {true|false}]`` - Is Sequential
Splitting of Data Operations Enabled? [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies whether I/O, such as write, copy-offload and
punch-holes, are split sequentially, rather than being run in parallel on
the source and destination. Spliiting I/O sequentially will make the
system more robust, and less prone to I/O errors. Starting 9.11.1,
enabling this feature improves the performance when the workload is NAS
based and is metadata heavy. However, it will make the IO performance
slower for large file workloads like LUNs, databases, virtualization
```

```
containers, etc. The default value for the parameter `-are-data-ops-
sequentially-split` is ``_false_`` . The parameter `-are-data-ops-
sequentially-split` should only be used if frequent I/O timeout or
"OutOfSync" has happened. The parameter `-are-data-ops-sequentially-split`
requires an effective cluster version of ONTAP 9.6.0 or later on both the
source and destination clusters.

``[-sequential-split-op-timeout-secs <integer>]`` - Sequential Split Op
Timeout in Seconds [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the op timeout value used when the splitting mode
is sequential. This parameter is used only for Sync relationships.
Supported values are from 15 to 25 seconds.

``[-disable-fast-resync {true|false}]`` - Disable Fast Resync [.advaced-
priv]#(privilege: advanced)#::
Specifies whether fast resync for a SnapMirror Synchronous relationship is
disabled or not. The default value for this parameter is ``_false_`` .

``[-discard-configs <network>,...]`` - Configurations Not Replicated
During Identity Preserve Vserver DR::
Specifies the configuration to be dropped during replication. The
supported values are:
+
  * ``_network_`` - Drops network interfaces, routes, and kerberos
configuration.

+
This parameter is supported only for policies of type ``_async-mirror_``
and applicable only for identity-preserve Vserver SnapMirror
relationships.

``[-transfer-schedule-name <text>]`` - Transfer Schedule Name::
 This optional parameter specifies the schedule which is used to update
the SnapMirror relationships.

``[-throttle <throttleType>]`` - Throttle (KB/sec)::
This optional parameter limits the network bandwidth used for transfers.
It configures for the relationships the maximum rate (in Kbytes/sec) at
which data can be transferred. If no throttle is configured, by default
the SnapMirror relationships fully use the network bandwidth available.
You can also configure the relationships to fully use the network
bandwidth available by explicitly setting the throttle to ``_unlimited_``
or ``_0_`` . The minimum effective throttle value is four Kbytes/sec, so
if you specify a throttle value between ``_1_`` and ``_4_`` , it will be
treated as ``_4_`` . For FlexGroup volume relationships, the throttle
value is applied individually to each constituent relationship.
```

```
== Examples

The following example creates a SnapMirror policy named ``_TieredBackup_``
on a Vserver named ``_vs0.example.com_`` .
```

vs0.example.com::> snapmirror policy create -vserver vs0.example.com
-policy TieredBackup -type vault -tries 10 -restart never

```
== Related Links

* xref:{relative_path}snapmirror-update.html[snapmirror update]
* xref:{relative_path}snapmirror-resync.html[snapmirror resync]
* xref:{relative_path}snapmirror-policy-add-rule.html[snapmirror policy
add-rule]
* xref:{relative_path}snapmirror-initialize.html[snapmirror initialize]
* xref:{relative_path}snapmirror-quiesce.html[snapmirror quiesce]
* xref:{relative_path}snapmirror-resume.html[snapmirror resume]




[[IDbe21560a10cc0aafeb204327a84a6ee8]]
= snapmirror policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a SnapMirror policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror policy delete` command deletes a SnapMirror policy. A
policy that is to be deleted must not be associated with any SnapMirror
relationship. The built-in policies cannot be deleted.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver for the SnapMirror policy.

``-policy <sm_policy>`` - SnapMirror Policy Name::
Specifies the SnapMirror policy name.

== Examples

The following example deletes a SnapMirror policy named ``_TieredBackup_``
on Vserver ``_vs0.example.com_`` :
```

vs0.example.com::> snapmirror policy delete -vserver vs0.example.com
-policy TieredBackup

```
[[IDe47ee680ae02fe20b7496e6ba010975d]]
= snapmirror policy modify-rule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an existing rule in SnapMirror policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror policy modify-rule` command can be used to modify the
retention count, preserve setting, warning threshold count, schedule, and
prefix for a rule in a SnapMirror policy. Reducing the retention count or
disabling the preserve setting for a rule in a SnapMirror policy might
result in the deletion of snapshots on the vault destination when the next
transfer by the xref:{relative_path}snapmirror-update.html[snapmirror
update] command occurs or when the next scheduled snapshot creation on the
destination for the rule occurs. Modifying a rule to add a schedule will
enable creation of snapshots on the SnapMirror destination. snapshots on
the source that have a SnapMirror label matching this rule will not be
selected for transfer. Schedule and prefix can only be modified for rules
associated with SnapMirror policies of type ``_vault_`` or ``_mirror-
vault_`` . A SnapMirror policy with rules must have at least one rule
without a schedule.

NOTE: The rules in SnapMirror policies of type ``_async-mirror_`` cannot
be modified.



== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver for the SnapMirror policy.

``-policy <sm_policy>`` - SnapMirror Policy Name::
Specifies the SnapMirror policy name.
```

```
``-snapmirror-label <text>`` - Snapshot Copy Label::
This parameter specifies the rule that is to be modified in a SnapMirror
policy.

``[-keep <text>]`` - Snapshot Copy Retention Count::
Specifies the maximum number of snapshots that are retained on the
SnapMirror destination volume for a rule. The total number of snapshots
retained for all the rules in a policy cannot exceed 1019. For all the
rules in SnapMirror policies of type ``_async-mirror_`` , this parameter
must be set to value ``_1_`` .

``[-preserve {true|false}]`` - Snapshot Copy Preserve Enabled::
Specifies the behavior when the snapshot retention count is reached on the
SnapMirror vault destination for the rule. The default value is
``_false_`` , which means that the oldest snapshot will be deleted to make
room for new ones only if the number of snapshots has exceeded the
retention count specified in the "keep" parameter. When set to ``_true_``
, and when the snapshots have reached the retention count, then an
incremental SnapMirror vault update transfer will fail or if the rule has
a schedule, snapshots will no longer be created on the SnapMirror
destination. For all the rules in SnapMirror policies of type ``_async-
mirror_`` this parameter must be set to value ``_false_`` .

``[-warn <integer>]`` - Warning Threshold Count::
Specifies the warning threshold count for the rule. The default value is
``_0_`` . When set to a value greater than zero, an event is generated
after the number of snapshots (for the particular rule) retained on a
SnapMirror vault destination reaches the specified warn limit. The
preserve parameter for the rule must be ``_true_`` to set the warn
parameter to a value greater than zero.

``[-schedule <text>]`` - Snapshot Copy Creation Schedule::
This optional parameter specifies the name of the schedule associated with
a rule. This parameter is allowed only for rules associated with
SnapMirror policies of type ``_vault_`` or ``_mirror-vault_`` . When this
parameter is specified, snapshots are directly created on the SnapMirror
destination. The snapshots created will have the same content as the
latest snapshot already present on the SnapMirror destination. snapshots
on the source that have a SnapMirror label matching this rule will not be
selected for transfer. The default value is ``_-_`` .
+
NOTE: You define and name a schedule using the xref:{relative_path}job-
schedule-cron-create.html[job schedule cron create] command.
+
```

```
``[-prefix <text>]`` - Snapshot Copy Creation Prefix::
This optional parameter specifies the prefix for the snapshot name to be
created as per the schedule. If no value is specified, then the
``_snapmirror-label_`` will be used as the prefix. The prefix parameter
can only be specified for rules which have a schedule.

``[-retention-period {{<integer> seconds|minutes|hours|days|months|years}
| infinite}]`` - Snapshot Copy Retention Period::
This optional parameter specifies the duration for which the snapshots
associated with this rule will be locked. The parameter is specified as a
number followed by a suffix. The valid suffixes are ``_seconds_`` ,
``_minutes_`` , ``_hours_`` , ``_days_`` , ``_months_`` , and ``_years_``
. For example, a value of ``_6months_`` represents a retention period of 6
months. It can also be set to a special value of ``_infinite_`` for
infinite retention. The destination volume must have the ``_snapshot-
locking-enabled_`` parameter set to true for this field to take effect
else it will be ignored.

== Examples

The following example changes the retention count for ``_nightly_``
snapshots to ``_6_`` for a rule named ``_nightly_`` on a SnapMirror policy
named ``_TieredBackup_`` on Vserver ``_vs0.example.com_`` :
```

vs0.example.com::> snapmirror policy modify-rule -vserver vs0.example.com
-policy TieredBackup -snapmirror-label nightly -keep 6

```
== Related Links

* xref:{relative_path}snapmirror-update.html[snapmirror update]
* xref:{relative_path}job-schedule-cron-create.html[job schedule cron
create]




[[ID27b43ecf8a3d902746777daf1ba35948]]
= snapmirror policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a SnapMirror policy
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror policy modify` command can be used to modify the policy
attributes.

NOTE: Use the xref:{relative_path}snapmirror-policy-modify-
rule.html[snapmirror policy modify-rule] command to modify a rule in a
SnapMirror policy.



== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver for the SnapMirror policy.

``-policy <sm_policy>`` - SnapMirror Policy Name::
Specifies the SnapMirror policy name.

``[-comment <text>]`` - Comment::
Specifies a text comment for the SnapMirror policy. If the comment
contains spaces, it must be enclosed within quotes.

``[-tries <integer_or_unlimited>]`` - Tries Limit::
Determines the maximum number of times to attempt each manual or scheduled
transfer for a SnapMirror relationship. The value of this parameter must
be a positive integer or ``_unlimited_`` . The default value is ``_8_`` .

``[-transfer-priority {low|normal}]`` - Transfer Scheduling Priority::
Specifies the priority at which a transfer runs. The supported values are
``_normal_`` or ``_low_`` . The ``_normal_`` transfers are scheduled
before the ``_low_`` priority transfers. The default is ``_normal_`` .

``[-ignore-atime {true|false}]`` - Ignore File Access Time::
This parameter applies only to extended data protection (XDP)
relationships. It specifies whether incremental transfers will ignore
files which have only their access time changed. The supported values are
``_true_`` or ``_false_`` . The default is ``_false_`` .

``[-restart {always|never|default}]`` - Restart Behavior::
This parameter applies only to data protection relationships. It defines
the behavior of SnapMirror if an interrupted transfer exists. The
supported values are ``_always_`` , ``_never_`` , or ``_default_`` . If
the value is set to ``_always_`` , an interrupted SnapMirror transfer
```

```
always restarts provided it has a restart checkpoint and the conditions
are the same as they were before the transfer was interrupted. In
addition, a new SnapMirror snapshot is created which will then be
transferred. If the value is set to ``_never_`` , an interrupted
SnapMirror transfer will never restart, even if a restart checkpoint
exists. A new SnapMirror snapshot will still be created and transferred.
ONTAP version 8.2 will interpret a value of ``_default_`` as being the
same as ``_always_`` . Vault transfers will always resume based on a
restart checkpoint, provided the snapshot still exists on the source
volume.

``[-is-network-compression-enabled {true|false}]`` - Is Network
Compression Enabled::
Specifies whether network compression is enabled for transfers. The
supported values are ``_true_`` or ``_false_`` . The default is
``_false_`` .

``[-rpo <integer>]`` - Recovery Point Objective (seconds)::
Specifies the time for recovery point objective, in seconds. This
parameter is only supported for a policy of type ``_continuous_`` .

``[-always-replicate-snapshots {true|false}]`` - This prioritizes
replication of app-consistent snapshots over synchronous replication::
If this parameter is set to true, it specifies that SnapMirror Synchronous
relationships will lose the zero RPO protection upon failure in
replicating application created snapshots. The default value is false.

``[-common-snapshot-schedule <text>]`` - Common Snapshot Copy Creation
Schedule for SnapMirror Synchronous [.advaced-priv]#(privilege:
advanced)#::
Specifies the common snapshot creating schedule. This parameter is only
supported for SnapMirror Synchronous relationships.

``[-are-data-ops-sequentially-split {true|false}]`` - Is Sequential
Splitting of Data Operations Enabled? [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies whether I/O, such as write, copy-offload and
punch-holes, are split sequentially, rather than being run in parallel on
the source and destination. Spliiting the I/O sequentially will make the
system more robust, and less prone to I/O errors. Starting 9.11.1,
enabling this feature improves the performance when the workload is NAS
based and is metadata heavy. However, it will make the IO performance
slower for large file workloads like LUNs, databases, virtualization
containers, etc. The default value of parameter `-sequential-split-data-
ops` is ``_false_`` . The parameter `-are-data-ops-sequentially-split`
should only be used if frequent I/O timeout or "OutOfSync" has happened.
```

```
Changes made by the `snapmirror policy modify -sequential-split-data-ops`
command do not take effect until the next resync. Changes do not affect
resync or initialize operations that have started and have not finished
yet. The parameter `-are- data-ops-sequentially-split` requires an
effective cluster version of ONTAP 9.6.0 or later on both the source and
destination clusters.

``[-sequential-split-op-timeout-secs <integer>]`` - Sequential Split Op
Timeout in Seconds [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the op timeout value used when the splitting mode
is sequential. This parameter is used only for Sync relationships.
Supported values are from 15 to 25 seconds. The default value of the
parameter `-sequential-split-op-timeout` is ``_15 seconds_`` . Changes
made by the `snapmirror policy modify -sequential-split-op-timeout-secs`
command do not take effect until the next resync. Changes do not affect
resync or initialize operations that have started and have not finished
yet.

``[-disable-fast-resync {true|false}]`` - Disable Fast Resync [.advaced-
priv]#(privilege: advanced)#::
Specifies whether fast resync for a SnapMirror Synchronous relationship is
disabled or not. The default value for this parameter when a policy is
created is ``_false_`` .

``[-discard-configs <network>,...]`` - Configurations Not Replicated
During Identity Preserve Vserver DR::
Specifies the configuration to be dropped during replication. The
supported values are:
+
  * ``_network_`` - Drops network interfaces, routes, and kerberos
configuration.

+
This parameter is supported only for policies of type ``_async-mirror_``
and applicable only for identity-preserve Vserver SnapMirror
relationships.

``[-transfer-schedule-name <text>]`` - Transfer Schedule Name::
 This optional parameter specifies the schedule which is used to update
the SnapMirror relationships.

``[-throttle <throttleType>]`` - Throttle (KB/sec)::
This optional parameter limits the network bandwidth used for transfers.
It configures for the relationships the maximum rate (in Kbytes/sec) at
which data can be transferred. If no throttle is configured, by default
the SnapMirror relationships fully use the network bandwidth available.
```

```
You can also configure the relationships to fully use the network
bandwidth available by explicitly setting the throttle to ``_unlimited_``
or ``_0_`` . The minimum effective throttle value is four Kbytes/sec, so
if you specify a throttle value between ``_1_`` and ``_4_`` , it will be
treated as ``_4_`` . For FlexGroup volume relationships, the throttle
value is applied individually to each constituent relationship.

== Examples

The following example changes the "transfer-priority" and the "comment"
text of a SnapMirror policy named ``_TieredBackup_`` on Vserver
``_vs0.example.com_`` :
```

vs0.example.com::> snapmirror policy modify -vserver vs0.example.com -policy TieredBackup -transfer-priority low -comment "Use for tiered backups"

```
== Related Links

* xref:{relative_path}snapmirror-policy-modify-rule.html[snapmirror policy
modify-rule]




[[ID9fdf05ce1a45cc6e4f5745b6a256c10c]]
= snapmirror policy remove-rule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a rule from SnapMirror policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror policy remove-rule` command removes a rule from a
SnapMirror policy. On the destination of a SnapMirror relationship with
`snapmirror policy` of type ``_vault_`` or ``_mirror-vault_`` , all
snapshots with a SnapMirror label matching the rule being removed are no
longer processed by SnapMirror and might need to be deleted manually. A
`snapmirror policy` of type ``_vault_`` must have at least one rule if
that policy is associated with a SnapMirror relationship. A SnapMirror
```

```
policy with rules must have at least one rule without a schedule.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver for the SnapMirror policy.

``-policy <sm_policy>`` - SnapMirror Policy Name::
Specifies the SnapMirror policy name.

``-snapmirror-label <text>`` - Snapshot Copy Label::
This parameter specifies the rule that is removed from the SnapMirror
policy.
+
The rule for SnapMirror label ``_sm_created_`` cannot be removed from
SnapMirror policies of type ``_async-mirror_`` or ``_mirror-vault_`` .

== Examples

The following example removes a rule named ``_nightly_`` from a SnapMirror
policy named ``_TieredBackup_`` on Vserver ``_vs0.example.com_`` :
```

vs0.example.com::> snapmirror policy remove-rule -vserver vs0.example.com -policy TieredBackup -snapmirror-label nightly

```
[[IDbfffe3badd896e762d8d52e74ca44683]]
= snapmirror policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show SnapMirror policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror policy show` command displays the following information
about SnapMirror policies:

* Vserver Name
* SnapMirror Policy Name
* SnapMirror Policy Type
```

```
* Number of Rules in the policy
* Tries
* Transfer Priority
* Comment for the policy
* Individual Rule Names
* Keep value for the Rule
* Total of Keep values across all Rules in the policy



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
Selects the policies that match this parameter value.

``[-policy <sm_policy>]`` - SnapMirror Policy Name::
Selects the policies that match this parameter value.

``[-type {vault|async-mirror|mirror-vault|strict-sync-mirror|sync-
mirror|automated-failover|automated-failover-duplex|continuous}]`` -
Snapmirror Policy Type::
Selects the policies that match this parameter value. A policy can be of
type ``_async-mirror_`` , ``_vault_`` or ``_mirror-vault_`` .

``[-owner {cluster-admin|vserver-admin}]`` - Owner of the Policy::
Selects the policies that match this parameter value. A policy can be
owned by either the ``_"Cluster Admin"_`` or a ``_"Vserver Admin"_`` .

``[-comment <text>]`` - Comment::
Selects the policies that match this parameter value.

``[-tries <integer_or_unlimited>]`` - Tries Limit::
Selects the policies that match this parameter value.

``[-transfer-priority {low|normal}]`` - Transfer Scheduling Priority::
Selects the policies that match this parameter value.

``[-ignore-atime {true|false}]`` - Ignore File Access Time::
```

```
Selects the policies that match this parameter value.

``[-restart {always|never|default}]`` - Restart Behavior::
Selects the policies that match this parameter value.

``[-is-network-compression-enabled {true|false}]`` - Is Network
Compression Enabled::
Selects the policies that match this parameter value.

``[-create-snapshot {true|false}]`` - Create a New Snapshot Copy::
Selects the policies that match this parameter value.

``[-rpo <integer>]`` - Recovery Point Objective (seconds)::
Selects the policies that match this parameter value.

``[-always-replicate-snapshots {true|false}]`` - This prioritizes
replication of app-consistent snapshots over synchronous replication::
Selects the policies that match this parameter value.

``[-snapmirror-label <text>,...]`` - Snapshot Copy Label::
Selects the policies that match this parameter value.

``[-keep <text>,...]`` - Snapshot Copy Retention Count::
Selects the policies that match this parameter value.

``[-preserve {true|false}]`` - Snapshot Copy Preserve Enabled::
Selects the policies that match this parameter value.

``[-warn <integer>,...]`` - Warning Threshold Count::
Selects the policies that match this parameter value.

``[-schedule <text>,...]`` - Snapshot Copy Creation Schedule::
Selects the policies that match this parameter value.

``[-prefix <text>,...]`` - Snapshot Copy Creation Prefix::
Selects the policies that match this parameter value.

``[-retention-period {{<integer> seconds|minutes|hours|days|months|years}
| infinite}]`` - Snapshot Copy Retention Period::
This optional parameter specifies the duration for which the snapshots
associated with this rule will be locked. The parameter is specified as a
number followed by a suffix. The valid suffixes are ``_seconds_`` ,
``_minutes_`` , ``_hours_`` , ``_days_`` , ``_months_`` , and ``_years_``
. For example, a value of ``_6months_`` represents a retention period of 6
months. It can also be set to a special value of ``_infinite_`` for
infinite retention. The destination volume must have the ``_snapshot-
locking-enabled_`` parameter set to true for this field to take effect
```

```
else it will be ignored.

``[-total-rules <integer>]`` - Total Rules in the Policy::
Selects the policies that match this parameter value.

``[-total-keep <integer>]`` - Total Retention Count for All Rules in the
Policy::
Selects the policies that match this parameter value.

``[-common-snapshot-schedule <text>]`` - Common Snapshot Copy Creation
Schedule for SnapMirror Synchronous [.advaced-priv]#(privilege:
advanced)#::
Selects the policies that match this parameter value.

``[-are-data-ops-sequentially-split {true|false}]`` - Is Sequential
Splitting of Data Operations Enabled? [.advaced-priv]#(privilege:
advanced)#::
Selects the policies that match this parameter value.

``[-sequential-split-op-timeout-secs <integer>]`` - Sequential Split Op
Timeout in Seconds [.advaced-priv]#(privilege: advanced)#::
Selects the policies that match this parameter value.

``[-disable-fast-resync {true|false}]`` - Disable Fast Resync [.advaced-
priv]#(privilege: advanced)#::
Selects the policies that match this parameter value.

``[-discard-configs <network>,...]`` - Configurations Not Replicated
During Identity Preserve Vserver DR::
Selects the policies that match this parameter value.

``[-transfer-schedule-name <text>]`` - Transfer Schedule Name::
Selects the policies that match this parameter value.

``[-throttle <throttleType>]`` - Throttle (KB/sec)::
Selects the policies that match this parameter value.

== Examples

The following example displays information about all SnapMirror policies:
```

cs::> snapmirror policy show
Vserver Policy Policy Number Transfer
Name Name Type Of Rules Tries Priority Comment
------- ------------------ ------ -------- ----- -------- ----------
cs DPDefault async-mirror 2 8 normal Asynchronous SnapMirror policy for mirroring all snapshots and the latest active file system.

SnapMirror Label: sm_created Keep: 1
all_source_snapshots 1
Total Keep: 2

cs MirrorAllSnapshots async-mirror 2 8 normal Asynchronous SnapMirror policy for mirroring all snapshots and the latest active file system.
SnapMirror Label: sm_created Keep: 1
all_source_snapshots 1
Total Keep: 2

cs MirrorAndVault mirror-vault 3 8 normal A unified Asynchronous SnapMirror and SnapVault policy for mirroring the latest active file system and daily and weekly snapshots.
SnapMirror Label: sm_created Keep: 1
daily 7
weekly 52
Total Keep: 60

cs MirrorLatest async-mirror 1 8 normal Asynchronous SnapMirror policy for mirroring the latest active file system.
SnapMirror Label: sm_created Keep: 1
Total Keep: 1

vs0.example.com
TieredBackup vault 0 8 normal Use for tiered backups
Snapmirror-label: - Keep: -
Total Keep: 0

cs Unified7year mirror-vault 4 8 normal Unified SnapMirror policy with 7year retention.
SnapMirror Label: sm_created Keep: 1
daily 7
weekly 52
monthly 84
Total Keep: 144

cs XDPDefault vault 2 8 normal Vault policy with daily and weekly rules.
SnapMirror Label: daily Keep: 7
weekly 52
Total Keep: 59

7 entries were displayed.

```
The following example shows all the policies with the following fields -
vserver (default), policy (default) and transfer-priority:
```

cs::> snapmirror policy show -fields transfer-priority
vserver policy transfer-priority
---------- --------- ----------
cs DPDefault normal
cs MirrorAllSnapshots
normal
cs MirrorAndVault
normal
cs MirrorLatest

normal
vs0.example.com
TieredBackup
normal
cs Unified7year
normal
cs XDPDefault
normal
7 entries were displayed.

```
:leveloffset: -1


= snapmirror snapshot-owner commands

:leveloffset: +1


[[IDb40421ca2b073a616903f1de8f0a0176]]
= snapmirror snapshot-owner create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add an owner to preserve a Snapshot copy for a SnapMirror mirror-to-vault
cascade configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror snapshot-owner create` command adds an owner on the
specified snapshot. A snapshot can have at most one owner. An owner can
only be added to a snapshot on a read-write volume. The snapshot must have
a valid SnapMirror label.

NOTE: Refer to the ``_ONTAP Data Protection Guide_`` for valid use cases
to add an owner on a snapshot.



== Parameters

``-vserver <vserver name>`` - Vserver Name::
```

```
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume.

``-snapshot <snapshot name>`` - Snapshot Copy Name::
This parameter specifies the name of the snapshot.

``[-owner <owner name>]`` - Snapshot Copy Owner Name::
This parameter specifies the name of the owner for the snapshot. The owner
name can be made up of the characters A to Z, a to z, 0 to 9, and "_". The
name can be up to 32 characters in length. When not specified, an owner
will be added with a system-generated default name.

== Examples

The following example adds owner ``_app1_`` on snapshot ``_snap1_`` on
volume ``_vol1_`` in Vserver ``_vs0.example.com_`` .
```

cluster1::> snapmirror snapshot-owner create -vserver vs0.example.com
-volume vol1 -snapshot snap1 -owner app1

```
The following example adds a default owner on snapshot ``_snap2_`` on
volume ``_vol1_`` in Vserver ``_vs0.example.com_`` .
```

cluster1::> snapmirror snapshot-owner create -vserver vs0.example.com
-volume vol1 -snapshot snap2

```
[[IDb1aa046bc6c87649ce1531d87d0a3b5f]]
= snapmirror snapshot-owner delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an owner used to preserve a Snapshot copy for a SnapMirror mirror-
to-vault cascade configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror snapshot-owner delete` command removes an owner on the
specified snapshot, which was added using the
xref:{relative_path}snapmirror-snapshot-owner-create.html[snapmirror
snapshot-owner create] command.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume.

``-snapshot <snapshot name>`` - Snapshot Copy Name::
This parameter specifies the name of the snapshot.

``[-owner <owner name>]`` - Snapshot Copy Owner Name::
This parameter specifies the name of the owner for the snapshot. When not
specified, the owner with the system-generated default name will be
removed.

== Examples

The following example removes owner ``_app1_`` on snapshot ``_snap1_`` on
volume ``_vol1_`` in Vserver ``_vs0.example.com_`` .
```

cluster1::> snapmirror snapshot-owner delete -vserver vs0.example.com
-volume vol1 -snapshot snap1 -owner app1

```
The following example removes the default owner on snapshot ``_snap2_`` on
volume ``_vol1_`` in Vserver ``_vs0.example.com_`` .
```

cluster1::> snapmirror snapshot-owner delete -vserver vs0.example.com
-volume vol1 -snapshot snap2

```
== Related Links

* xref:{relative_path}snapmirror-snapshot-owner-create.html[snapmirror
snapshot-owner create]




[[IDca376fe5da47e59fc4d550b7f39939db]]
= snapmirror snapshot-owner show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Snapshot Copies with Owners

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `snapmirror snapshot-owner show` command is used to list all snapshots
with owners that were added using the xref:{relative_path}snapmirror-
snapshot-owner-create.html[snapmirror snapshot-owner create] command.

== Parameters

{ ``[-fields <fieldname>,...]``::
If this parameter is specified, the command displays information about the
specified fields.

| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed information
about all fields.

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.
```

```
``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume.

``[-snapshot <snapshot name>]`` - Snapshot Copy Name::
If this parameter is specified, the command displays the owner name for
the specified snapshot.

== Examples

The following example lists all snapshots with owners on volume ``_vol1_``
in Vserver ``_vs0_`` . The system-generated default owner name is
displayed as "-".
```

cluster1::> snapmirror snapshot-owner show
-vserver vs0.example.com -volume vol1
Vserver Volume Snapshot Owner Names
-------- ---------- ---------- ----------
vs0.example.com
vol1 snap2 -
snap1 app1

```
The following example displays the owner name for snapshot ``_snap1_`` on
volume ``_vol1_`` in Vserver ``_vs0.example.com_`` .
```

cluster1::> snapmirror snapshot-owner show
-vserver vs0.example.com -volume vol1 -snapshot snap1
Vserver: vs0.example.com
Volume: vol1
Snapshot: snap1
Owner Names: app1

```
== Related Links

* xref:{relative_path}snapmirror-snapshot-owner-create.html[snapmirror
snapshot-owner create]




:leveloffset: -1


:leveloffset: -1
```

```
= statistics commands

:leveloffset: +1


[[IDfd463f06e89bac99d6a374ca8afe86e6]]
= statistics show-periodic
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Continuously display current performance data at regular interval

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command continuously displays specified performance data at a regular
interval. The command output displays data in the following columns:

NOTE: This command has been deprecated and may be removed from a future
version of ONTAP. Use the "statistics show" command with the tabular
format instead.



* cpu avg: Average processor utilization across all processors in the
system.
* cpu busy: Overall system utilization based on CPU utilization and
subsystem utilization. Examples of subsystems include the storage
subsystem and RAID subsystem.
* total ops: Number of total operations per second.
* nfs-ops: Number of NFS operations per second.
* cifs-ops: Number of CIFS operations per second.
* fcache ops: Number of 7M FlexCache operations per second.
* pkts recv: Number of packets received over physical ports per second.
* pkts sent: Number of packets sent over physical ports per second.
* total recv: Total network traffic received over physical ports per
second (KBps).
* total sent: Total network traffic sent over physical ports per second
(KBps).
* data busy: The percentage of time that data ports sent or received data.
* data recv: Network traffic received on data ports (KBps).
* data sent: Network traffic sent on data ports (KBps).
```

```
* cluster busy: The percentage of time that cluster ports sent or received
data.
* cluster recv: Network traffic received on cluster ports (KBps).
* cluster sent: Network traffic sent on cluster ports (KBps).
* disk read: Data read from disk (KBps).
* disk write: Data written to disk (KBps).



== Parameters

``[-object <text>]`` - Object [.advaced-priv]#(privilege: advanced)#::
Selects the object for which you want to display performance data. The
default object is "``_cluster_`` ".

``[-instance <text>]`` - Instance [.advaced-priv]#(privilege: advanced)#::
Selects the instance for which you want to display performance data. This
parameter is required if you specify the `-object` parameter and enter any
object other than "``_cluster_`` ". Multiple values for this parameter are
not supported.
+
For example, if you want to display disk object statistics, you can use
this parameter to specify the name of a specific disk whose statistics you
want to view.

``[-counter <text>]`` - Counter [.advaced-priv]#(privilege: advanced)#::
Selects the counters for which you want to display performance data. If
you do not specify this parameter, the command displays statistics for all
of the counters in the specified objects. To specify multiple counters,
use "|" between each counter.

``[-preset <text>]`` - Preset [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics for the
specified preset.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes for which you want to display performance data. The
default node is "``_cluster:summary_`` ".

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Selects the Vserver for which you want to display performance data. If you
do not specify this parameter, the command displays statistics for all of
the Vservers in the cluster.

``[-interval <integer>]`` - Interval in Seconds [.advaced-
```

```
priv]#(privilege: advanced)#::
Specifies, in seconds, the interval between statistics updates. The
default setting is 1 second.

``[-iterations <integer>]`` - Number of Iterations [.advaced-
priv]#(privilege: advanced)#::
Specifies the number of iterations the command runs before terminating.
The default setting is 0 (zero); this means that the command continues to
run until you interrupt it by pressing Ctrl-C.

``[-summary {true|false}]`` - Print Summary [.advaced-priv]#(privilege:
advanced)#::
Specifies whether the command prints a final summary of statistics after
the command has gone through all of its iterations. The default setting is
`true` .

``[-filter <text>]`` - Filter Data [.advaced-priv]#(privilege:
advanced)#::
Selects instances that match the specified filter criteria. For example,
to display instances from node1, specify `-filter "node_name=node1"` .

== Examples

The following example displays the "cluster" statistics for a node named
node1. Because no number of iterations is specified, this command will
continue to run until you interrupt it by pressing Ctrl-C.
```

```
cluster1::*> statistics show-periodic -node node1
cpu cpu cpu total fcache pkts pkts total total data data data cluster cluster cluster disk disk
avg busy total ops nfs-ops cifs-ops ops recv sent recv sent busy recv sent busy recv sent read write
---- ---- ----- -------- -------- -------- -------- -------- -------- -------- -------- ---- -------- -------- ------- -------- -------- --------
--------
6% 6% 26% 0 0 0 0 641 39 126KB 4.68KB 0% 7.33KB 0B 0% 111KB 4.68KB 0B 0B
5% 5% 21% 0 0 0 0 254 7 16.7KB 1.75KB 0% 5.60KB 0B 0% 5.48KB 1.75KB 0B 0B
6% 6% 24% 0 0 0 0 307 41 21.2KB 5.32KB 0% 6.13KB 0B 0% 8.99KB 5.32KB 0B 0B
4% 4% 17% 0 0 0 0 364 16 23.8KB 2.58KB 0% 8.88KB 0B 0% 6.05KB 2.58KB 0B 0B
10% 10% 42% 0 0 0 0 673 7 124KB 1.92KB 0% 9.82KB 0B 0% 104KB 1.92KB 0B 0B
7% 7% 28% 0 0 0 0 407 38 28.1KB 4.38KB 0% 8.79KB 0B 0% 10.5KB 4.38KB 106KB 528KB
4% 4% 19% 0 0 0 0 328 16 21.6KB 2.58KB 0% 7.27KB 0B 0% 7.02KB 2.58KB 0B 0B
5% 5% 22% 0 0 0 0 324 31 21.9KB 4.35KB 0% 6.99KB 0B 0% 7.95KB 4.35KB 0B 0B
5% 5% 21% 0 0 0 0 242 16 16.1KB 2.60KB 0% 5.10KB 0B 0% 5.89KB 2.60KB 0B 0B
4% 4% 17% 0 0 0 0 273 16 18.0KB 2.60KB 0% 5.91KB 0B 0% 6.20KB 2.60KB 0B 0B
```

```
The following example displays the "processor" statistics for an instance
named processor1. This command will display only five iterations.
```

```
cluster1::*> statistics show-periodic -object processor -instance processor1 -iteration 5
```

instance processor sk
name busy switches
-------- --------- --------
processor1 8% 1722
processor1 6% 1234
processor1 5% 1680
processor1 4% 1336
processor1 7% 1801

```
The following example displays the processor statistics for an instance
named processor1 and counters "processor_busy" and "sk_switches". This
command will display only five iterations.
```

cluster1::*> statistics show-periodic -object processor -instance processor1 -iteration 5 -counter processor_busy|sk_switches
processor sk
busy switches
--------- --------
5% 1267
4% 1163
7% 1512
5% 1245
4% 1128

```
[[ID986f144922cee6e4bc48af9378c72498]]
= statistics show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display performance data for a time interval

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command displays performance data for a period of time .

To display data for a period of time, collect a sample using the
xref:{relative_path}statistics-start.html[statistics start] and
xref:{relative_path}statistics-stop.html[statistics stop] commands. The
data that displays is calculated data based on the samples the cluster
collects. To view the sample, specify the -sample-id parameter.

== 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.

| ``[-tab  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is specified, the command displays performance data in
tabular format.

``[-object <text>]`` - Object [.advaced-priv]#(privilege: advanced)#::
Selects the objects for which you want to display performance data. To
view a list of valid object names, type `statistics show``-object```_?_``
or xref:{relative_path}statistics-catalog-object-show.html[statistics
catalog object show] . To specify multiple objects, use "|" between each
object.
+
NOTE: You should limit the scope of this command to only a few objects at
a time to avoid a potentially significant impact on the performance of the
system.
+


``[-instance <text>]`` - Instance [.advaced-priv]#(privilege: advanced)#::
Selects the instances for which you want to display performance data. If
you do not specify this parameter, the command displays statistics for all
of the instances associated with the specified objects. To specify
multiple instances, use "|" between each instance.
+
For example, if you want to display disk object statistics, you can use
this parameter to specify the name of a specific disk whose statistics you
want to view. If you do not specify this parameter, the command displays
statistics for all disks in the system.

``[-counter <text>]`` - Counter [.advaced-priv]#(privilege: advanced)#::
Selects the counters for which you want to display performance data. To
specify multiple counters, use "|" between each counter.

``[-preset <text>]`` - Preset [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics for the
specified preset.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes for which you want to display performance data.
```

```
``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Selects the Vserver for which you want to display performance data.

``[-value <text>]`` - Text Value [.advaced-priv]#(privilege: advanced)#::
Selects the performance data that matches the specified counter value.

``[-labels <text>,...]`` - List of Labels [.advaced-priv]#(privilege:
advanced)#::
Selects the performance data that matches the specified label.

``[-values <text>,...]`` - List of Values [.advaced-priv]#(privilege:
advanced)#::
Displays only the statistics that have the specified values.

``[-filter <text>]`` - Filter Data [.advaced-priv]#(privilege:
advanced)#::
Selects performance data for the instance that matches the specified
filter criteria. For example, to display the instances that match a value
of greater than 50 for the total_ops counter, specify `-filter
"total_ops>50"` .

``[-sample-id <text>]`` - Sample Identifier [.advaced-priv]#(privilege:
advanced)#::
Displays performance data for the specified sample. You collect a sample
by using the statistics start and statistics stop commands.

``[-interval <integer>]`` - Interval [.advaced-priv]#(privilege:
advanced)#::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``[-iterations <integer>]`` - Iterations [.advaced-priv]#(privilege:
advanced)#::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``[-sort-key <text>]`` - Counter Used For Sorting [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics sorted by
the specified counter. Only one counter can be specified.

``[-sort-order {ascending|descending}]`` - Sort Order [.advaced-
priv]#(privilege: advanced)#::
This parameter may be used in conjunction with the `-sort-key` parameter.
This parameter changes the order in which statistics are sorted. Possible
```

```
values are ``_ascending_`` and ``_descending_`` . The default setting is
``_descending_`` .

``[-max <integer>]`` - Tracker Size [.advaced-priv]#(privilege:
advanced)#::
Specifies the number of most active instances of an active object to
display. The default setting is to display all of the instances.

== Examples

The following example starts collecting statistics and displays statistics
for the sample named smpl_1 for counters: avg_processor_busy and cpu_busy
```

cluster1::*> statistics start -object system -counter avg_processor_busy|cpu_busy -sample-id smpl_1
Statistics collection is being started for Sample-id: smpl_1

cluster1::*> statistics show -sample-id smpl_1
Object: system
Instance: cluster
Start-time: 8/2/2012 18:27:53
End-time: 8/2/2012 18:27:56
Cluster: cluster1
Counter Value
--- ---
avg_processor_busy 6%
cpu_busy 6%

```
The following example starts and stops data collection and displays
statistics for the sample named smpl_1 for counters: avg_processor_busy
and cpu_busy
```

cluster1::*> statistics start -object system -counter avg_processor_busy|cpu_busy -sample-id smpl_1
Statistics collection is being started for Sample-id: smpl_1

cluster1::*> statistics stop -sample-id smpl_1
Statistics collection is being stopped for Sample-id: smpl_1

cluster1::*> statistics show -sample-id smpl_1
Object: system
Instance: cluster
Start-time: 8/2/2012 18:27:53
End-time: 8/2/2012 18:27:56
Cluster: cluster1
Counter Value
--- ---
avg_processor_busy 6%
cpu_busy 6%

```
== Related Links

* xref:{relative_path}statistics-start.html[statistics start]
* xref:{relative_path}statistics-stop.html[statistics stop]
* xref:{relative_path}statistics-catalog-object-show.html[statistics
catalog object show]




[[IDc384997b33b63be3e5c4f5654e0f6c42]]
= statistics start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start data collection for a sample

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command starts the collection of performance data. Use the
xref:{relative_path}statistics-stop.html[statistics stop] command to stop
the collection. You view the sample of performance data by using the
xref:{relative_path}statistics-show.html[statistics show] command. You can
collect more than one sample at a time.

== Parameters

``[-object <text>]`` - Object [.advaced-priv]#(privilege: advanced)#::
Selects the objects for which you want to collect performance data. This
parameter is required. To view a list of valid object names, type
xref:{relative_path}statistics-catalog-object-show.html[statistics catalog
object show] at the command prompt. To specify multiple objects, use "|"
between each object.
+
NOTE: You should limit the scope of this command to only a few objects at
a time to avoid a potentially significant impact on the performance of the
system.
+


``[-instance <text>]`` - Instance [.advaced-priv]#(privilege: advanced)#::
```

```
Selects the instances for which you want to collect performance data. If
you do not specify this parameter, the command collects statistics for all
of the instances associated with the specified objects. To specify
multiple instances, use "|" between each instance.
+
For example, if you want to collect disk object statistics, you can use
this parameter to specify the name of a specific disk whose statistics you
want to view. If you do not specify this parameter, the command will
collect statistics for all disks in the system.

``[-counter <text>]`` - Counter [.advaced-priv]#(privilege: advanced)#::
Selects the counters for which you want to collect performance data. If
you do not specify this parameter, the command collects statistics for all
of the counters in the specified objects. To specify multiple counters,
use "|" between each counter.

``[-preset <text>]`` - Preset [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics for the
specified preset.

``[-sample-id <text>]`` - Sample Identifier [.advaced-priv]#(privilege:
advanced)#::
Specifies an identifier for the sample. Identifiers must be unique and are
restricted to the characters 0-9, a-z, A-Z, and "_". If you do not specify
this parameter, the command generates a sample identifier for you and
defines this sample as the default sample for the CLI session. When you
run the xref:{relative_path}statistics-show.html[statistics show] command
without specifying the `-sample-id` parameter, data from the default
sample displays. If you run this command during the same CLI session and
do not specify the `-sample-id` parameter, the command overwrites the
previous sample. The command does not delete the default sample when you
close your session.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Selects the vserver for which you want to collect performance data. If you
do not specify this parameter, the command collects statistics for all of
the Vservers in the cluster.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects the node for which you want to collect performance data. If you do
not specify this parameter, the command collects statistics for all of the
nodes in the cluster.

``[-filter <text>]`` - Filter [.advaced-priv]#(privilege: advanced)#::
```

```
Selects performance data for the instance that matches the specified
filter criteria. For example, to display the instances from node1, specify
`-filter "node_name=node1"` .

``[-duration <integer>]`` - Sample Duration in Minutes [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command will collect the closing
sample after the time specified. Duration can be specified in minutes.

``[-max <integer>]`` - Tracker Size [.advaced-priv]#(privilege:
advanced)#::
Specifies the number of most active instances of an active object to
display. The default setting is to display all of the instances.

``[-sort-key <text>]`` - Counter Used For Sorting [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics sorted by
the specified counter. Only one counter can be specified.

``[-sort-order {ascending|descending}]`` - Sort Order [.advaced-
priv]#(privilege: advanced)#::
This parameter may be used in conjunction with the -sort-key parameter.
This parameter changes the order in which statistics are sorted. Possible
values are ascending and descending. The default setting is descending.

== Examples

The following example starts statistics collection for sample "smpl_1":
```

cluster1::*> statistics start -object system -sample-id smpl_1
Statistics collection is being started for Sample-id: smpl_1

```
The following example starts collecting statistics for the sample named
smpl_1 for counters: avg_processor_busy and cpu_busy
```

cluster1::*> statistics start -object system -counter avg_processor_busy|cpu_busy -sample-id smpl_1
Statistics collection is being started for Sample-id: smpl_1

```
== Related Links

* xref:{relative_path}statistics-stop.html[statistics stop]
* xref:{relative_path}statistics-show.html[statistics show]
* xref:{relative_path}statistics-catalog-object-show.html[statistics
catalog object show]




[[ID4edc7fc19f983d4f499cc8ffd888a845]]
= statistics stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop data collection for a sample

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command stops the collection of performance data. You view the sample
of performance data by using the xref:{relative_path}statistics-
show.html[statistics show] command.

== Parameters

``[-sample-id <text>]`` - Sample Identifier [.advaced-priv]#(privilege:
advanced)#::
Specifies the identifier of the sample for which you want to stop data
collection. If you do not specify this parameter, the command stops data
collection for the last sample that you started by running the
xref:{relative_path}statistics-start.html[statistics start] command
without the `-sample-id` parameter.

== Examples

The following example stops data collection for sample "smpl_1":
```

cluster1::*> statistics stop -sample-id smpl_1
Statistics collection is being stopped for Sample-id: smpl_1

```
== Related Links

* xref:{relative_path}statistics-show.html[statistics show]
* xref:{relative_path}statistics-start.html[statistics start]




= statistics aggregate commands

:leveloffset: +1


[[ID4a5ebb3315be93683bd1f32c5893eb1d]]
= statistics aggregate show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Aggregate throughput and latency metrics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continuously displays performance data for aggregates at a
regular interval. The command output displays data in the following
columns:

* Aggregate - aggregate name.
* Node - node name.
* Total Ops - total number of operations per second.
* Read Ops - read operations per second.
* Write Ops - write operations per second.



== Parameters

``[-aggregate <text>]`` - Aggregate::
Selects the aggregate for which you want to display performance data.

``[-node {<nodename>|local}]`` - Node::
Selects the node for which you want to display performance data.
```

```
``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies maximum number of aggregates to display. The default setting is
25.

== Examples

The following example displays aggregate statistics:
```

cluster1::> statistics aggregate show
cluster-1 : 12/31/1969 16:00:04

```
                                     *Total Read Write
            Aggregate          Node    Ops  Ops   Ops
--------------------- ------------- ------ ---- -----
aggr0_cluster_node2_0 cluster-node2      9    0     8
                 aggr0 cluster-node1      6    0     5
```

```
:leveloffset: -1


= statistics cache commands

:leveloffset: +1


[[ID08c6eb0a8da946213363695a096c2c17]]
= statistics cache flash-pool show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Flash pool throughput metrics

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command continuously displays performance data for Flash Pool caches
at a regular interval. The command output displays data in the following
columns:

* Aggregate - aggregate name.
* Vserver - vserver name.
* Volume - volume name.
* Read Hit - percent of IOs serviced from a cache level.
* Write Hit - percent of IOs serviced from a cache level.
* Cache Used - SSD cache blocks used.
* Read Blocks - read blocks.
* Write Blocks - write blocks.
* Rejects - cache rejects.



== Parameters

``[-aggregate <text>]`` - Aggregate::
Selects the aggregate for which you want to display performance data.

``[-vserver <vserver name>]`` - Vserver::
Selects the vserver for which you want to display performance data.

``[-volume <text>]`` - Volume::
Selects the volume for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
```

```
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of Flash Pool aggregates to display. The
default setting is 25.

== Examples

The following example displays Flash Pool statistics:
```

cluster1::> statistics cache flash-pool show
cluster1 : 12/31/2013 16:00:04
Read Write
Hit Hit Cache Read Write
Aggregate Vserver Volume (%) (%) Used Blocks Blocks Rejects
--- --- --- --- --- --- --- --- ---
aggr1 - -total- 0 0 0 0 0 0
aggr2 vs1 vol1 0 0 0 0 0 0

```
:leveloffset: -1


= statistics catalog commands

:leveloffset: +1


[[ID2a7469f3afe561a8a93dc59a62bcf40e]]
= statistics catalog counter show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of counters in an object

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command displays the names and descriptions of counters. The
displayed data is either node-specific or cluster-wide, depending on the
objects specified.
```

```
== 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.

| ``[-describe  ]`` [.advaced-priv]#(privilege: advanced)# }::
Displays detailed information about each counter, including privilege
level, label, and whether the counter is a key counter.

``-object <text>`` - Object [.advaced-priv]#(privilege: advanced)#::
Selects the object for which you want to display the list of counters.
This parameter is required. To view a list of valid object names, type
`statistics catalog counter show``-object```_?_`` or
xref:{relative_path}statistics-catalog-object-show.html[statistics catalog
object show] .

``[-counter <text>]`` - Counter [.advaced-priv]#(privilege: advanced)#::
Selects the counters that match this parameter value. If you do not
specify this parameter, the command displays details for all counters.

``[-filter <text>]`` - Filter Data [.advaced-priv]#(privilege:
advanced)#::
Selects the counters that match this parameter value. For example, to
display counters from node1, specify `-filter "node_name=node1"` .

``[-label <text>,...]`` - Labels for Array Counters [.advaced-
priv]#(privilege: advanced)#::
Selects the counters that match this parameter value. A label is the name
of the bucket to which an array counter belongs.

``[-description <text>]`` - Description [.advaced-priv]#(privilege:
advanced)#::
Selects the counters that match this parameter value.

``[-privilege <text>]`` - Privilegel Level [.advaced-priv]#(privilege:
advanced)#::
Selects the counters that match this parameter value.

``[-is-key-counter {true|false}]`` - Is Key Counter [.advaced-
priv]#(privilege: advanced)#::
Selects the counters that are key counters (true) or are not key counters
(false). A key counter is a counter that uniquely identifies an instance
across the cluster. The default setting is false. For example,
"vserver_name" and "node_name" are key counters because they identify the
```

```
specific Vserver or node to which the instance belongs.

``[-is-deprecated {true|false}]`` - Is Counter Deprecated [.advaced-
priv]#(privilege: advanced)#::
Selects the counters that are deprecated (true) or are not deprecated
(false).

``[-replaced-by <text>]`` - Replaced By Counter If Deprecated [.advaced-
priv]#(privilege: advanced)#::
Selects all deprecated counters that are replaced by the counter provided
to this parameter.

== Examples

The following example displays the list of counters in the processor
object.
```

cluster1::> statistics catalog counter show -object processor
Object: processor
Counter Description
--- ---
instance_name Instance Name
instance_uuid Instance UUID
node_name System node name
node_uuid System node id
process_name Ontap process that provided this instance
processor_busy Percentage of elapsed time that the processor
is executing non-idle processes
processor_elapsed_time Wall-clock time since boot used for
calculating processor utilization
sk_switches Number of sk switches per second
8 entries were displayed.

```
== Related Links

* xref:{relative_path}statistics-catalog-object-show.html[statistics
catalog object show]




[[ID578cfd5440ab200dda33fc37e9c85e59]]
= statistics catalog instance show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display the list of instances associated with an object

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command displays the names of instances associated with the specified
object. The displayed data is either node-specific or cluster-wide,
depending on the objects specified.

== 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.

``-object <text>`` - Object [.advaced-priv]#(privilege: advanced)#::
Selects the object for which you want to display the list of instances.
This parameter is required. To view a list of valid object names, type
`statistics catalog instance show``-object```_?_`` or
xref:{relative_path}statistics-catalog-object-show.html[statistics catalog
object show] .

``[-instance <text>]`` - Instance Name [.advaced-priv]#(privilege:
advanced)#::
Selects the instances that match this parameter value. If you do not
specify this parameter, the command displays all the instances.

``[-filter <text>]`` - Filter Data [.advaced-priv]#(privilege:
advanced)#::
Selects the instances that match this parameter value. For example, to
display instances from vserver1, specify `-filter "vserver_name=vserver1"`
.

``[-vserver <vserver name>,...]`` - Vserver Name [.advaced-
priv]#(privilege: advanced)#::
Selects the instances that match this parameter value. If you do not
specify this parameter, the command displays instances for all of the
Vservers in the cluster.

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Selects the instances that match this parameter value. If you do not
specify this parameter, the command displays instances for all of the
```

```
nodes in the cluster.

== Examples

The following example displays the list of instances associated with the
processor object.
```

cluster1::> statistics catalog instance show -object processor
Object: processor
processor0
processor0
processor1
processor1
4 entries were displayed.

```
== Related Links

* xref:{relative_path}statistics-catalog-object-show.html[statistics
catalog object show]




[[ID538f0497c5561db08418b825f5d56294]]
= statistics catalog object show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of objects

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command displays the names and descriptions of objects from which you
can obtain performance data. The displayed data is either node-specific or
cluster-wide, depending on the objects specified.

== 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.

| ``[-describe  ]`` [.advaced-priv]#(privilege: advanced)# }::
Displays detailed information about each object, including privilege
level.

``[-object <text>]`` - Object [.advaced-priv]#(privilege: advanced)#::
Selects the objects for which you want to display information. If you do
not specify this parameter, the command displays details for all of the
objects.

``[-privilege <text>]`` - Privilege Level [.advaced-priv]#(privilege:
advanced)#::
Selects the objects that match this parameter value.

``[-is-deprecated {true|false}]`` - Is Object Deprecated [.advaced-
priv]#(privilege: advanced)#::
Selects the objects that are deprecated (true) or are not deprecated
(false).

``[-replaced-by <text>]`` - Replaced By Object If Deprecated [.advaced-
priv]#(privilege: advanced)#::
Selects all deprecated objects that are replaced by the object provided to
this parameter.

``[-is-statistically-tracked {true|false}]`` - Is Object Statistically
Tracked [.advaced-priv]#(privilege: advanced)#::
Specifies if the object is statistically tracked

``[-description <text>]`` - Description [.advaced-priv]#(privilege:
advanced)#::
Selects the objects that match this parameter value.

== Examples

The following example displays descriptions of all objects in the cluster:
```

cluster1::> statistics catalog object show
aggregate CM object for exporting aggregate performance
counters
audit_ng CM object for exporting audit_ng performance
counters
cifs These counters report activity from both SMB
and SMB2 revisions of the CIFS protocol. For
information isolated to SMB, see the 'smb1'
object. For SMB2, see the 'smb2' object.
cifs:node These counters report activity from both SMB

and SMB2 revisions of the CIFS protocol. For information isolated to SMB, see the 'smb1' object. For SMB2, see the 'smb2' object.
cifs:vserver These counters report activity from both SMB and SMB2 revisions of the CIFS protocol. For information isolated to SMB, see the 'smb1' object. For SMB2, see the 'smb2' object.
cluster_peer The cluster peer object contains peer counters.

```
:leveloffset: -1


= statistics disk commands

:leveloffset: +1


[[IDfdd11f3553438234a7bd52ba066df584]]
= statistics disk show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disk throughput and latency metrics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command continuously displays performance data for disks at a regular
interval. The command output displays data in the following columns:

* Disk - disk name.
* Node - node name.
* Busy (%) - percentage of time there was at least one outstanding request
to the disk.
* Total Ops - total operations per second.
* Read Ops - read operations per second.
* Write Ops - write operations per second.



== Parameters
```

```
``[-disk <text>]`` - Disk [.advaced-priv]#(privilege: advanced)#::
Selects the disk for which you want to display performance data.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects the node for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval [.advaced-priv]#(privilege:
advanced)#::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations [.advaced-priv]#(privilege:
advanced)#::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances [.advaced-
priv]#(privilege: advanced)#::
Specifies the maximum number of disks to display. The default setting is
25.

== Examples

The following example displays disk statistics:
```

cluster1::> statistics disk show
cluster1 : 12/31/1969 16:00:04

```
                   Busy *Total Read Write
Disk         Node   (%)    Ops  Ops   Ops
-------- -------- ---- ------ ---- -----
VMw-1.31    node2    0      2    2     0
VMw-1.30    node2    0      3    0     3
VMw-1.3     node1    0      0    0     0
VMw-1.29    node2    0      1    0     1
```

```
:leveloffset: -1
```

```
= statistics lif commands

:leveloffset: +1


[[ID2cc02e336bc1200c5ce82f66b3f34186]]
= statistics lif show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Logical network interface throughput and latency metrics

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command continuously displays performance data for LIFs at a regular
interval. The command output displays data in the following columns:

* LIF - logical interface name.
* Vserver - vserver name.
* Recv Packet - packets received per second.
* Recv Data (Bps) - bytes received per second.
* Recv Errors - receive errors per second.
* Sent Packet - packets sent per second.
* Sent Data (Bps) - bytes sent per second.
* Sent Errors - transfer errors per second.
* Current Port - current use port.



== Parameters

``[-lif <text>]`` - LIF::
Selects the LIF for which you want to display performance data.

``[-vserver <vserver name>]`` - Vserver::
Selects the vserver for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
```

```
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of LIFs to display. The default setting is
25.

== Examples

The following example displays LIFs statistics:
```

cluster1::> statistics lif show
cluster1 : 12/31/1969 16:00:04
Recv Sent
Recv Data Recv Sent Data Sent Current
LIF Vserver Packet (Bps) Errors Packet (Bps) Errors Port
---------- ---------- ------ ----- ------ ------ ----- ------ -------
node2_clus_1 Cluster 3 536 0 3 338 0 e0a
node2_clus_2 Cluster 3 398 0 3 287 0 e0b
node1_clus_2 Cluster 3 338 0 3 536 0 e0b
node1_clus_1 Cluster 3 287 0 3 398 0 e0a
node2_mgmt1 ncluster-1 0 0 0 0 0 0 e0c
node1_mgmt1 ncluster-1 0 0 0 0 0 0 e0c

```
:leveloffset: -1


= statistics lun commands

:leveloffset: +1


[[ID48e873e2440e3d88cad544391c5f701f]]
= statistics lun show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
LUN throughput and latency metrics

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command continuously displays performance data for LUNs at a regular
interval. The command output displays data in the following columns:

* Lun - LUN name.
* Vserver - vserver name.
* Total Ops - total number of operations per second.
* Read Ops - read operations per second.
* Write Ops - write operations per second.
* Other Ops - other operations per second.
* Read (Bps) - read throughput in bytes per second.
* Write (Bps) - write throughput in bytes per second.
* Latency(us) - average latency for an operation in microseconds.



== Parameters

``[-lun <text>]`` - Lun::
Selects the LUN for which you want to display performance data.

``[-vserver <vserver name>]`` - Vserver::
Selects the vserver for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of LUNs to display. The default setting is
25.
```

```
== Examples

The following example displays LUN statistics:
```

cluster1::> statistics lun show
cluster1 : 12/31/2013 16:00:04
*Total Read Write Other Read Write Latency
Lun Vserver Ops Ops Ops Ops (Bps) (Bps) (us)
---- ------- ------ ---- ----- ----- ------ ----- -------
lun1 vs1 58 13 15 29 310585 3014 39
lun0 vs2 56 0 11 45 8192 28826 47

```
:leveloffset: -1


= statistics namespace commands

:leveloffset: +1


[[IDb010f0b7a11035739b9bc04d851d86ee]]
= statistics namespace show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Namespace throughput and latency metrics

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command continuously displays performance data for Namespaces at a
regular interval. The command output displays data in the following
columns:

* Namespace - Namespace name.
* Vserver - vserver name.
* Total Ops - total number of operations per second.
* Read Ops - read operations per second.
* Write Ops - write operations per second.
* Other Ops - other operations per second.
```

```
* Read (Bps) - read throughput in bytes per second.
* Write (Bps) - write throughput in bytes per second.
* Latency(ms) - average latency for an operation in milliseconds.



== Parameters

``[-namespace <text>]`` - Namespace::
Selects the Namespace for which you want to display performance data.

``[-vserver <vserver name>]`` - Vserver::
Selects the vserver for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of Namespaces to display. The default setting
is 25.

== Examples

The following example displays Namespace statistics:
```

cluster1::> statistics namespace show
cluster1 : 12/31/2017 16:00:04
*Total Read Write Other Read Write Latency
Namespace Vserver Ops Ops Ops Ops (Bps) (Bps) (ms)
--- --- --- --- --- --- --- --- ---
ns1 vs1 58 13 15 29 310585 3014 39
ns0 vs2 56 0 11 45 8192 28826 47

```
:leveloffset: -1



= statistics nfs commands
```

```
:leveloffset: +1


[[ID2c87b508a60dcb6b6ae7635945b21602]]
= statistics nfs show-mount
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display mount statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `statistics nfs show-mount` command displays the following statistics
about the NFS mounts on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of mount operations
* Total number of dump operations
* Total number of unmount operations
* Total number of unmountall operations
* Total number of export operations
* Total number of exportall operations
* Total number of pathconf operations
* Total number of all the above operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only for
the specified node.

``[-result {success|failure|all}]`` - Result [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-mount <Counter with Delta>]`` - Mount Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of mount operations.

``[-dump <Counter with Delta>]`` - Dump Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of dump operations.

``[-unmnt <Counter with Delta>]`` - UnMount Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmount operations.

``[-unmntall <Counter with Delta>]`` - UnMountAll Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmountall operations.

``[-export <Counter with Delta>]`` - Export Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of export operations.

``[-exportall <Counter with Delta>]`` - ExportAll Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of exportall operations.

``[-pathconf <Counter with Delta>]`` - PathConf Operations [.advaced-
priv]#(privilege: advanced)#::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of pathconf operations.

``[-total <Counter64 with Delta>]`` - Total Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total operations.

== Examples

The following example displays statistics about the NFS mounts for a node
named node1:
```

```
cluster1::*> statistics nfs show-mount -node node1

Node Value Delta
node1 ----------success-------
Null Ops: 2 0/s:16s
Mount Ops: 1 0/s:16s
Dump Ops: 0 -
Unmount Ops: 1 0/s:16s
Unmount All Ops: 0 -
Export Ops: 0 -
ExportAll Ops 0 -
PathConf Ops: 0 -
Total Ops: 4 0/s:16s

Node Value Delta
node1 ----------failure-------
Null Ops: 0 -
Mount Ops: 0 -
Dump Ops: 0 -
Unmount Ops: 0 -
Unmount All Ops: 0 -
Export Ops: 0 -
ExportAll Ops 0 -
PathConf Ops: 0 -
Total Ops: 0 -
```

```
[[ID72bfcf770aaa213985ff80c10871c14d]]
= statistics nfs show-nlm
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Display NLM statistics
```

```
*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `statistics nfs show-nlm` command displays the following statistics
about the Network Lock Manager (NLM) on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of test operations
* Total number of lock operations
* Total number of cancel operations
* Total number of unlock operations
* Total number of granted operations
* Total number of share operations
* Total number of unshare operations
* Total number of nmlock operations
* Total number of freeall operations
* Total number of all the above operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

NOTE: This command requires an effective cluster version earlier than
ONTAP 9.0. Data for nodes running ONTAP 9.0 or later is not collected, and
will not be displayed. Use the xref:{relative_path}statistics-
show.html[statistics show]`-object```_nlm_`` command instead.



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only for
the specified node.
```

```
``[-result {success|failure|all}]`` - Result [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-test <Counter with Delta>]`` - Test Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of test operations.

``[-lock <Counter with Delta>]`` - Lock Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lock operations.

``[-cancel <Counter with Delta>]`` - Cancel Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of cancel operations.

``[-unlock <Counter with Delta>]`` - Unlock Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unlock operations.

``[-granted <Counter with Delta>]`` - Granted Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of granted operations.

``[-share <Counter with Delta>]`` - Share Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of share operations.

``[-unshare <Counter with Delta>]`` - Unshare Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unshare operations.

``[-nmlock <Counter with Delta>]`` - NmLock Operations [.advaced-
priv]#(privilege: advanced)#::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of nmlock operations.

``[-freeall <Counter with Delta>]`` - FreeAll Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of freeall operations.

``[-total <Counter64 with Delta>]`` - Total Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total operations.

== Examples

The following example displays statistics about the NLM for a node named
node1:
```

cluster1::*> statistics nfs show-nlm -node node1

Node Value Delta
node1 --------success-------
Null: 0 -
Test: 0 -
Lock: 2 0/s:23s
Cancel: 0 -
Unlock: 1 0/s:23s
Granted: 0 -
Share: 0 -
Unshare: 0 -
NmLock: 0 -
FreeAll: 0 -
Total: 3 0/s:23s

Node Value Delta
node1 --------failure-------
Null: 0 -
Test: 0 -
Lock: 0 -
Cancel: 0 -
Unlock: 0 -
Granted: 0 -
Share: 0 -
Unshare: 0 -
NmLock: 0 -
FreeAll: 0 -
Total: 0 -

```
== Related Links
```

```
* xref:{relative_path}statistics-show.html[statistics show]




[[ID4c58f3adcd93f3a8fc1a73bad48fa0c9]]
= statistics nfs show-statusmon
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status monitor statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `statistics nfs show-statusmon` command displays the following
statistics about the Status Monitor on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of stat operations
* Total number of monitor operations
* Total number of unmonitor operations
* Total number of unmonitor all operations
* Total number of simucrash operations
* Total number of notify operations
* Total number of all the above operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only for
the specified node.

``[-result {success|failure|all}]`` - Result [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-stat <Counter with Delta>]`` - Stat Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of stat operations.

``[-monitor <Counter with Delta>]`` - Monitor Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of monitor operations.

``[-unmonitor <Counter with Delta>]`` - Unmonitor Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmonitor operations.

``[-unmonall <Counter with Delta>]`` - Unmonitor All Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmonitor all
operations.

``[-simucrash <Counter with Delta>]`` - SimuCrash Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of simucrash operations.

``[-notify <Counter with Delta>]`` - Notify Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of notify operations.

``[-total <Counter64 with Delta>]`` - Total Operations [.advaced-
```

```
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total operations.

== Examples

The following example displays statistics about the status monitor for a
node named node1:
```

cluster1::*> statistics nfs show-statusmon -node node1

Node Value Delta
node1 --------success-------
Null Ops: 0 -
Stat Ops: 0 -
Monitor Ops: 0 -
Unmonitor Ops: 0 -
Unmon All Ops: 0 -
SimuCrash Ops: 0 -
Notify Ops: 0 -
Total Ops: 0 -

Node Value Delta
node1 --------failure-------
Null Ops: 0 -
Stat Ops: 0 -
Monitor Ops: 0 -
Unmonitor Ops: 0 -
Unmon All Ops: 0 -
SimuCrash Ops: 0 -
Notify Ops: 0 -
Total Ops: 0 -

```
[[IDe130ae257e8691e0293e90862b4d92c3]]
= statistics nfs show-v3
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NFSv3 statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `statistics nfs show-v3` command displays the following statistics
```

```
about the NFSv3 operations on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of getattr operations
* Total number of setattr operations
* Total number of lookup operations
* Total number of access operations
* Total number of readsymlink operations
* Total number of read operations
* Total number of write operations
* Total number of create operations
* Total number of mkdir operations
* Total number of symlink operations
* Total number of mknod operations
* Total number of remove operations
* Total number of rmdir operations
* Total number of rename operations
* Total number of link operations
* Total number of readdir operations
* Total number of readdirplus operations
* Total number of fsstat operations
* Total number of fsinfo operations
* Total number of pathconf operations
* Total number of commit operations
* Total number of nfsv3 operations
* Percent of null operations
* Percent of getattr operations
* Percent of setattr operations
* Percent of lookup operations
* Percent of access operations
* Percent of readsymlink operations
* Percent of read operations
* Percent of write operations
* Percent of create operations
* Percent of mkdir operations
* Percent of symlink operations
* Percent of mknod operations
* Percent of remove operations
* Percent of rmdir operations
* Percent of rename operations
* Percent of link operations
* Percent of readdir operations
* Percent of readdirplus operations
* Percent of fsstat operations
* Percent of fsinfo operations
```

```
* Percent of pathconf operations
* Percent of commit operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays NFSv3 statistics only
for the specified node.

``[-result {success|failure|all}]`` - Result::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-gattr <Counter with Delta>]`` - GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getattr operations.

``[-sattr <Counter with Delta>]`` - SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setattr operations.

``[-lookup <Counter with Delta>]`` - LookUp Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookup operations.

``[-access <Counter with Delta>]`` - Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of access operations.

``[-rsym <Counter with Delta>]`` - ReadSymlink Operations::
If you specify this parameter, the command displays statistics only about
```

```
the node or nodes that have the specified number of readsymlink
operations.

``[-read <Counter with Delta>]`` - Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of read operations.

``[-write <Counter with Delta>]`` - Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of write operations.

``[-create <Counter with Delta>]`` - Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of create operations.

``[-mkdir <Counter with Delta>]`` - MkDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of mkdir operations.

``[-symln <Counter with Delta>]`` - SymLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of symlink operations.

``[-mknod <Counter with Delta>]`` - MkNod Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of mknod operations.

``[-remove <Counter with Delta>]`` - Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of remove operations.

``[-rmdir <Counter with Delta>]`` - RmDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rmdir operations.

``[-rename <Counter with Delta>]`` - Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rename operations.

``[-link <Counter with Delta>]`` - Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of link operations.

``[-rdir <Counter with Delta>]`` - ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readdir operations.
```

```
``[-rdirp <Counter with Delta>]`` - ReadDirPlus Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readdirplus
operations.

``[-fsstat <Counter with Delta>]`` - FsStat Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of fsstat operations.

``[-fsinfo <Counter with Delta>]`` - FsInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of fsinfo operations.

``[-pconf <Counter with Delta>]`` - PathConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of pathconf operations.

``[-commit <Counter with Delta>]`` - Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of commit operations.

``[-total <Counter64 with Delta>]`` - Total Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total NFSv3
operations.

``[-null-pct <Counter with Delta>]`` - Percent Null Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of null operations.

``[-gattr-pct <Counter with Delta>]`` - Percent GetAttr Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getattr
operations.

``[-sattr-pct <Counter with Delta>]`` - Percent SetAttr Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setattr
operations.

``[-lookup-pct <Counter with Delta>]`` - Percent LookUp Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookup operations.

``[-access-pct <Counter with Delta>]`` - Percent Access Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of access operations.
```

```
``[-rsym-pct <Counter with Delta>]`` - Percent ReadSymlink Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readsymlink
operations.

``[-read-pct <Counter with Delta>]`` - Percent Read Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of read operations.

``[-write-pct <Counter with Delta>]`` - Percent Write Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of write operations.

``[-create-pct <Counter with Delta>]`` - Percent Create Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of create operations.

``[-mkdir-pct <Counter with Delta>]`` - Percent MkDir Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of mkdir operations.

``[-symln-pct <Counter with Delta>]`` - Percent SymLink Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of symlink
operations.

``[-mknod-pct <Counter with Delta>]`` - Percent MkNod Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of mknod operations.

``[-remove-pct <Counter with Delta>]`` - Percent Remove Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of remove operations.

``[-rmdir-pct <Counter with Delta>]`` - Percent RmDir Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rmdir operations.

``[-rename-pct <Counter with Delta>]`` - Percent Rename Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rename operations.

``[-link-pct <Counter with Delta>]`` - Percent Link Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of link operations.
```

```
``[-rdir-pct <Counter with Delta>]`` - Percent ReadDir Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdir
operations.

``[-rdirp-pct <Counter with Delta>]`` - Percent ReadDirPlus Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdirplus
operations.

``[-fsstat-pct <Counter with Delta>]`` - Percent FsStat Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of fsstat operations.

``[-fsinfo-pct <Counter with Delta>]`` - Percent FsInfo Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of fsinfo operations.

``[-pconf-pct <Counter with Delta>]`` - Percent PathConf Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of pathconf
operations.

``[-commit-pct <Counter with Delta>]`` - Percent Commit Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of commit operations.

== Examples

The following example displays statistics about the NFSv3 operations for a
node named node1:
```

cluster1::> statistics nfs show-v3 -node node1

Node Value Delta Percent Ops Delta
node1 ----------success----------
Null Ops: 4 - 7% -
GetAttr Ops: 10 - 19% -
SetAttr Ops: 2 - 4% -
Lookup Ops: 2 - 4% -
Access Ops: 14 - 26% -
ReadSymlink Ops: 0 - 0% -
Read Ops: 0 - 0% -
Write Ops: 0 - 0% -
Create Ops: 2 - 4% -
MkDir Ops: 1 - 2% -
Symlink Ops: 0 - 0% -
MkNod Ops: 0 - 0% -

Remove Ops: 1 - 2% -
RmDir Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Link Ops: 0 - 0% -
ReadDir Ops: 2 - 4% -
ReadDirPlus Ops: 10 - 19% -
FsStat Ops: 1 - 2% -
FsInfo Ops: 5 - 9% -
PathConf Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Total Ops: 54 -

Node Value Delta Percent Ops Delta
node1 ----------failure----------
Null Ops: 0 - 0% -
GetAttr Ops: 0 - 0% -
SetAttr Ops: 0 - 0% -
Lookup Ops: 2 - 100% -
Access Ops: 0 - 0% -
ReadSymlink Ops: 0 - 0% -
Read Ops: 0 - 0% -
Write Ops: 0 - 0% -
Create Ops: 0 - 0% -
MkDir Ops: 0 - 0% -
Symlink Ops: 0 - 0% -
MkNod Ops: 0 - 0% -
Remove Ops: 0 - 0% -
RmDir Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Link Ops: 0 - 0% -
ReadDir Ops: 0 - 0% -
ReadDirPlus Ops: 0 - 0% -
FsStat Ops: 0 - 0% -
FsInfo Ops: 0 - 0% -
PathConf Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Total Ops: 2 -

Node Value Delta Percent Ops Delta
node1 ---------- all ----------
Null Ops: 4 - 7% -
GetAttr Ops: 10 - 18% -
SetAttr Ops: 2 - 4% -
Lookup Ops: 4 - 7% -
Access Ops: 14 - 25% -
ReadSymlink Ops: 0 - 0% -
Read Ops: 0 - 0% -
Write Ops: 0 - 0% -
Create Ops: 2 - 4% -
MkDir Ops: 1 - 2% -
Symlink Ops: 0 - 0% -
MkNod Ops: 0 - 0% -
Remove Ops: 1 - 2% -
RmDir Ops: 0 - 0% -

Rename Ops: 0 - 0% -
Link Ops: 0 - 0% -
ReadDir Ops: 2 - 4% -
ReadDirPlus Ops: 10 - 18% -
FsStat Ops: 1 - 2% -
FsInfo Ops: 5 - 9% -
PathConf Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Total Ops: 56 -

```
[[ID851b1de28aa85987a9f49a1c438a4e92]]
= statistics nfs show-v4
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NFSv4 statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `statistics nfs show-v4` command displays the following statistics
about the NFSv4 operations on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of compound operations
* Total number of access operations
* Total number of close operations
* Total number of commit operations
* Total number of create operations
* Total number of delegpurge operations
* Total number of delegret operations
* Total number of getattr operations
* Total number of getfh operations
* Total number of link operations
* Total number of lock operations
* Total number of lockt operations
* Total number of locku operations
* Total number of lookup operations
* Total number of lookupp operations
* Total number of nverify operations
* Total number of open operations
```

```
* Total number of openattr operations
* Total number of openconf operations
* Total number of opendowng operations
* Total number of putfh operations
* Total number of putpubfh operations
* Total number of putrootfh operations
* Total number of read operations
* Total number of readdir operations
* Total number of readlink operations
* Total number of remove operations
* Total number of rename operations
* Total number of renew operations
* Total number of restorefh operations
* Total number of savefh operations
* Total number of secinfo operations
* Total number of setattr operations
* Total number of setcliid operations
* Total number of setcliidconf operations
* Total number of verify operations
* Total number of write operations
* Total number of rellockown operations
* Total number of total operations
* Percent of null operations
* Percent of compound operations
* Percent of access operations
* Percent of close operations
* Percent of commit operations
* Percent of create operations
* Percent of delegpurge operations
* Percent of delegret operations
* Percent of getattr operations
* Percent of getfh operations
* Percent of link operations
* Percent of lock operations
* Percent of lockt operations
* Percent of locku operations
* Percent of lookup operations
* Percent of lookupp operations
* Percent of nverify operations
* Percent of open operations
* Percent of openattr operations
* Percent of openconf operations
* Percent of opendowng operations
* Percent of putfh operations
* Percent of putpubfh operations
* Percent of putrootfh operations
```

```
* Percent of read operations
* Percent of readdir operations
* Percent of readlink operations
* Percent of remove operations
* Percent of rename operations
* Percent of renew operations
* Percent of restorefh operations
* Percent of savefh operations
* Percent of secinfo operations
* Percent of setattr operations
* Percent of setcliid operations
* Percent of setCliidconf operations
* Percent of verify operations
* Percent of write operations
* Percent of rellockown operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays NFSv4 statistics only
for the specified node.

``[-result {success|failure|all}]`` - Result::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-cmpnd <Counter with Delta>]`` - Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of compound operations.
```

```
``[-access <Counter with Delta>]`` - Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of access operations.

``[-close <Counter with Delta>]`` - Close Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of close operations.

``[-commit <Counter with Delta>]`` - Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of commit operations.

``[-create <Counter with Delta>]`` - Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of create operations.

``[-delpur <Counter with Delta>]`` - Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegpurge operations.

``[-delrtn <Counter with Delta>]`` - Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegret operations.

``[-gattr <Counter with Delta>]`` - GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getattr operations.

``[-getfh <Counter with Delta>]`` - GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getfh operations.

``[-link <Counter with Delta>]`` - Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of link operations.

``[-lock <Counter with Delta>]`` - Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lock operations.

``[-lockt <Counter with Delta>]`` - LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lockt operations.

``[-locku <Counter with Delta>]`` - LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of locku operations.
```

```
``[-lookup <Counter with Delta>]`` - Lookup Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookup operations.

``[-lookpp <Counter with Delta>]`` - LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookupp operations.

``[-nverfy <Counter with Delta>]`` - Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of nverify operations.

``[-open <Counter with Delta>]`` - Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of open operations.

``[-opattr <Counter with Delta>]`` - OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openattr operations.

``[-opconf <Counter with Delta>]`` - OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openconf operations.

``[-opndg <Counter with Delta>]`` - OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of opendowng operations.

``[-putfh <Counter with Delta>]`` - PutFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putfh operations.

``[-putpfh <Counter with Delta>]`` - PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putpubfh operations.

``[-putrfh <Counter with Delta>]`` - PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putrootfh operations.

``[-read <Counter with Delta>]`` - Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of read operations.

``[-readdr <Counter with Delta>]`` - ReadDir Operations::
If you specify this parameter, the command displays statistics only about
```

```
the node or nodes that have the specified number of readdir operations.

``[-rlink <Counter with Delta>]`` - ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readlink operations.

``[-remove <Counter with Delta>]`` - Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of remove operations.

``[-rename <Counter with Delta>]`` - Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rename operations.

``[-renew <Counter with Delta>]`` - Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of renew operations.

``[-restfh <Counter with Delta>]`` - RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of restorefh operations.

``[-savefh <Counter with Delta>]`` - SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of savefh operations.

``[-secinf <Counter with Delta>]`` - SecInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of secinfo operations.

``[-sattr <Counter with Delta>]`` - SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setattr operations.

``[-sclid <Counter with Delta>]`` - SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliid operations.

``[-scidc <Counter with Delta>]`` - SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliidconf
operations.

``[-verify <Counter with Delta>]`` - Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of verify operations.
```

```
``[-write <Counter with Delta>]`` - Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of write operations.

``[-relown <Counter with Delta>]`` - RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rellockown operations.

``[-total <Counter64 with Delta>]`` - Total Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total nfsv4
operations.

``[-null-pct <Counter with Delta>]`` - Percent Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of null operations.

``[-cmpnd-pct <Counter with Delta>]`` - Percent Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of compound
operations.

``[-access-pct <Counter with Delta>]`` - Percent Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of access operations.

``[-close-pct <Counter with Delta>]`` - Percent Close Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of close operations.

``[-commit-pct <Counter with Delta>]`` - Percent Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of commit operations.

``[-create-pct <Counter with Delta>]`` - Percent Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of create operations.

``[-delpur-pct <Counter with Delta>]`` - Percent Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegpurge
operations.

``[-delrtn-pct <Counter with Delta>]`` - Percent Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegret
operations.
```

```
``[-gattr-pct <Counter with Delta>]`` - Percent GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getattr
operations.

``[-getfh-pct <Counter with Delta>]`` - Percent GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getfh operations.

``[-link-pct <Counter with Delta>]`` - Percent Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of link operations.

``[-lock-pct <Counter with Delta>]`` - Percent Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lock operations.

``[-lockt-pct <Counter with Delta>]`` - Percent LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lockt operations.

``[-locku-pct <Counter with Delta>]`` - Percent LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of locku operations.

``[-lookup-pct <Counter with Delta>]`` - Percent Lookup Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookup operations.

``[-lookpp-pct <Counter with Delta>]`` - Percent LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookupp
operations.

``[-nverfy-pct <Counter with Delta>]`` - Percent Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of nverify
operations.

``[-open-pct <Counter with Delta>]`` - Percent Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of open operations.

``[-opattr-pct <Counter with Delta>]`` - Percent OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openattr
```

```
operations.

``[-opconf-pct <Counter with Delta>]`` - Percent OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openconf
operations.

``[-opndg-pct <Counter with Delta>]`` - Percent OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of opendowng
operations.

``[-putfh-pct <Counter with Delta>]`` - Percent PutFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putfh operations.

``[-putpfh-pct <Counter with Delta>]`` - Percent PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putpubfh
operations.

``[-putrfh-pct <Counter with Delta>]`` - Percent PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putrootfh
operations.

``[-read-pct <Counter with Delta>]`` - Percent Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of read operations.

``[-readdr-pct <Counter with Delta>]`` - Percent ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdir
operations.

``[-rlink-pct <Counter with Delta>]`` - Percent ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readlink
operations.

``[-remove-pct <Counter with Delta>]`` - Percent Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of remove operations.

``[-rename-pct <Counter with Delta>]`` - Percent Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rename operations.
```

```
``[-renew-pct <Counter with Delta>]`` - Percent Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of renew operations.

``[-restfh-pct <Counter with Delta>]`` - Percent RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of restorefh
operations.

``[-savefh-pct <Counter with Delta>]`` - Percent SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of savefh operations.

``[-secinf-pct <Counter with Delta>]`` - Percent SecInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of secinfo
operations.

``[-sattr-pct <Counter with Delta>]`` - Percent SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setattr
operations.

``[-sclid-pct <Counter with Delta>]`` - Percent SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliid
operations.

``[-scidc-pct <Counter with Delta>]`` - Percent SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliidconf
operations.

``[-verify-pct <Counter with Delta>]`` - Percent Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of verify operations.

``[-write-pct <Counter with Delta>]`` - Percent Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of write operations.

``[-relown-pct <Counter with Delta>]`` - Percent RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rellockown
operations.
```

```
== Examples

The following example displays statistics about the NFSv4 operations for a
node named node1:
```

cluster1::> statistics nfs show-v4 -node node1

Node Value Delta Percent Ops Delta
node1 ----------success----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 6% -
Close Ops: 8 - 3% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 27% -
Getfh Ops: 22 - 8% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 13 - 5% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 8 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 92 - 32% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -
Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 5% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 8 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 286 -
Node Value Delta Percent Ops Delta
node1 ----------failure----------
Null Procs: 0 - 0% -
Cmpnd Procs: 0 - -

Access Ops: 0 - 0% -
Close Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 0 - 0% -
Getfh Ops: 0 - 0% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 5 - 63% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 2 - 25% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 0 - 0% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 0 - 0% -
Read Ops: 0 - 0% -
Readdir Ops: 0 - 0% -
Readlink Ops: 0 - 0% -
Remove Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Renew Ops: 0 - 0% -
Restorefh Ops: 0 - 0% -
Savefh Ops: 0 - 0% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 1 - 13% -
Setclid Ops: 0 - 0% -
Setclidconf Ops: 0 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 0 - 0% -
Rlockown Ops: 0 - 0% -
Total Ops: 8 -
Node Value Delta Percent Ops Delta
node1 ---------- all ----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 5% -
Close Ops: 8 - 3% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 26% -
Getfh Ops: 22 - 7% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -

Lookup Ops: 18 - 6% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 10 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 92 - 31% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -
Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 4% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 9 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 294 -

```
[[IDea472b7a7e339aa778f57c5d40bea194]]
= statistics nfs show-v41
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NFSv4.1 statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `statistics nfs show-v41` command displays the following statistics
about the NFSv4.1 operations on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of compound operations
* Total number of access operations
* Total number of close operations
```

```
* Total number of commit operations
* Total number of create operations
* Total number of delegpurge operations
* Total number of delegret operations
* Total number of getattr operations
* Total number of getfh operations
* Total number of link operations
* Total number of lock operations
* Total number of lockt operations
* Total number of locku operations
* Total number of lookup operations
* Total number of lookupp operations
* Total number of nverify operations
* Total number of open operations
* Total number of openattr operations
* Total number of openconf operations
* Total number of opendowng operations
* Total number of putfh operations
* Total number of putpubfh operations
* Total number of putrootfh operations
* Total number of read operations
* Total number of readdir operations
* Total number of readlink operations
* Total number of remove operations
* Total number of rename operations
* Total number of renew operations
* Total number of restorefh operations
* Total number of savefh operations
* Total number of secinfo operations
* Total number of setattr operations
* Total number of setcliid operations
* Total number of setcliidconf operations
* Total number of verify operations
* Total number of write operations
* Total number of rellockown operations
* Total number of total operations
* Percent of null operations
* Percent of compound operations
* Percent of access operations
* Percent of close operations
* Percent of commit operations
* Percent of create operations
* Percent of delegpurge operations
* Percent of delegret operations
* Percent of getattr operations
* Percent of getfh operations
```

```
* Percent of link operations
* Percent of lock operations
* Percent of lockt operations
* Percent of locku operations
* Percent of lookup operations
* Percent of lookupp operations
* Percent of nverify operations
* Percent of open operations
* Percent of openattr operations
* Percent of openconf operations
* Percent of opendowng operations
* Percent of putfh operations
* Percent of putpubfh operations
* Percent of putrootfh operations
* Percent of read operations
* Percent of readdir operations
* Percent of readlink operations
* Percent of remove operations
* Percent of rename operations
* Percent of renew operations
* Percent of restorefh operations
* Percent of savefh operations
* Percent of secinfo operations
* Percent of setattr operations
* Percent of setcliid operations
* Percent of setCliidconf operations
* Percent of verify operations
* Percent of write operations
* Percent of rellockown operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays NFSv4.1 statistics
```

```
only for the specified node.

``[-result {success|failure|all}]`` - Result::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-cmpnd <Counter with Delta>]`` - Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of compound operations.

``[-access <Counter with Delta>]`` - Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of access operations.

``[-close <Counter with Delta>]`` - Close Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of close operations.

``[-commit <Counter with Delta>]`` - Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of commit operations.

``[-create <Counter with Delta>]`` - Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of create operations.

``[-delpur <Counter with Delta>]`` - Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegpurge operations.

``[-delrtn <Counter with Delta>]`` - Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegret operations.

``[-gattr <Counter with Delta>]`` - GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getattr operations.

``[-getfh <Counter with Delta>]`` - GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getfh operations.

``[-link <Counter with Delta>]`` - Link Operations::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of link operations.

``[-lock <Counter with Delta>]`` - Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lock operations.

``[-lockt <Counter with Delta>]`` - LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lockt operations.

``[-locku <Counter with Delta>]`` - LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of locku operations.

``[-lookup <Counter with Delta>]`` - Lookup Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookup operations.

``[-lookpp <Counter with Delta>]`` - LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookupp operations.

``[-nverfy <Counter with Delta>]`` - Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of nverify operations.

``[-open <Counter with Delta>]`` - Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of open operations.

``[-opattr <Counter with Delta>]`` - OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openattr operations.

``[-opconf <Counter with Delta>]`` - OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openconf operations.

``[-opndg <Counter with Delta>]`` - OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of opendowng operations.

``[-putfh <Counter with Delta>]`` - PutFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putfh operations.
```

```
``[-putpfh <Counter with Delta>]`` - PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putpubfh operations.

``[-putrfh <Counter with Delta>]`` - PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putrootfh operations.

``[-read <Counter with Delta>]`` - Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of read operations.

``[-readdr <Counter with Delta>]`` - ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readdir operations.

``[-rlink <Counter with Delta>]`` - ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readlink operations.

``[-remove <Counter with Delta>]`` - Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of remove operations.

``[-rename <Counter with Delta>]`` - Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rename operations.

``[-renew <Counter with Delta>]`` - Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of renew operations.

``[-restfh <Counter with Delta>]`` - RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of restorefh operations.

``[-savefh <Counter with Delta>]`` - SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of savefh operations.

``[-secinf <Counter with Delta>]`` - SecInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of secinfo operations.

``[-sattr <Counter with Delta>]`` - SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setattr operations.
```

```
``[-sclid <Counter with Delta>]`` - SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliid operations.

``[-scidc <Counter with Delta>]`` - SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliidconf
operations.

``[-verify <Counter with Delta>]`` - Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of verify operations.

``[-write <Counter with Delta>]`` - Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of write operations.

``[-relown <Counter with Delta>]`` - RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rellockown operations.

``[-total <Counter64 with Delta>]`` - Total Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total nfsv4.1
operations.

``[-null-pct <Counter with Delta>]`` - Percent Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of null operations.

``[-cmpnd-pct <Counter with Delta>]`` - Percent Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of compound
operations.

``[-access-pct <Counter with Delta>]`` - Percent Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of access operations.

``[-close-pct <Counter with Delta>]`` - Percent Close Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of close operations.

``[-commit-pct <Counter with Delta>]`` - Percent Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of commit operations.
```

```
``[-create-pct <Counter with Delta>]`` - Percent Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of create operations.

``[-delpur-pct <Counter with Delta>]`` - Percent Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegpurge
operations.

``[-delrtn-pct <Counter with Delta>]`` - Percent Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegret
operations.

``[-gattr-pct <Counter with Delta>]`` - Percent GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getattr
operations.

``[-getfh-pct <Counter with Delta>]`` - Percent GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getfh operations.

``[-link-pct <Counter with Delta>]`` - Percent Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of link operations.

``[-lock-pct <Counter with Delta>]`` - Percent Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lock operations.

``[-lockt-pct <Counter with Delta>]`` - Percent LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lockt operations.

``[-locku-pct <Counter with Delta>]`` - Percent LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of locku operations.

``[-lookup-pct <Counter with Delta>]`` - Percent Lookup Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookup operations.

``[-lookpp-pct <Counter with Delta>]`` - Percent LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookupp
```

```
operations.

``[-nverfy-pct <Counter with Delta>]`` - Percent Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of nverify
operations.

``[-open-pct <Counter with Delta>]`` - Percent Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of open operations.

``[-opattr-pct <Counter with Delta>]`` - Percent OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openattr
operations.

``[-opconf-pct <Counter with Delta>]`` - Percent OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openconf
operations.

``[-opndg-pct <Counter with Delta>]`` - Percent OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of opendowng
operations.

``[-putfh-pct <Counter with Delta>]`` - Percent PutFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putfh operations.

``[-putpfh-pct <Counter with Delta>]`` - Percent PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putpubfh
operations.

``[-putrfh-pct <Counter with Delta>]`` - Percent PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putrootfh
operations.

``[-read-pct <Counter with Delta>]`` - Percent Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of read operations.

``[-readdr-pct <Counter with Delta>]`` - Percent ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdir
```

```
operations.

``[-rlink-pct <Counter with Delta>]`` - Percent ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readlink
operations.

``[-remove-pct <Counter with Delta>]`` - Percent Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of remove operations.

``[-rename-pct <Counter with Delta>]`` - Percent Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rename operations.

``[-renew-pct <Counter with Delta>]`` - Percent Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of renew operations.

``[-restfh-pct <Counter with Delta>]`` - Percent RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of restorefh
operations.

``[-savefh-pct <Counter with Delta>]`` - Percent SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of savefh operations.

``[-secinf-pct <Counter with Delta>]`` - Percent SecInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of secinfo
operations.

``[-sattr-pct <Counter with Delta>]`` - Percent SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setattr
operations.

``[-sclid-pct <Counter with Delta>]`` - Percent SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliid
operations.

``[-scidc-pct <Counter with Delta>]`` - Percent SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliidconf
operations.
```

```
``[-verify-pct <Counter with Delta>]`` - Percent Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of verify operations.

``[-write-pct <Counter with Delta>]`` - Percent Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of write operations.

``[-relown-pct <Counter with Delta>]`` - Percent RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rellockown
operations.

== Examples

The following example displays statistics about the NFSv4.1 operations for
a node named node1:
```

cluster1::> statistics nfs show-v41 -node node1

Node Value Delta Percent Ops Delta
node1 ----------success----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 6% -
Close Ops: 8 - 3% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 27% -
Getfh Ops: 22 - 8% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 13 - 5% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 8 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 92 - 32% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -

Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 5% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 8 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 286 -
Node Value Delta Percent Ops Delta
node1 ----------failure----------
Null Procs: 0 - 0% -
Cmpnd Procs: 0 - -
Access Ops: 0 - 0% -
Close Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 0 - 0% -
Getfh Ops: 0 - 0% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 5 - 63% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 2 - 25% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 0 - 0% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 0 - 0% -
Read Ops: 0 - 0% -
Readdir Ops: 0 - 0% -
Readlink Ops: 0 - 0% -
Remove Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Renew Ops: 0 - 0% -
Restorefh Ops: 0 - 0% -
Savefh Ops: 0 - 0% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 1 - 13% -
Setclid Ops: 0 - 0% -
Setclidconf Ops: 0 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 0 - 0% -
Rlockown Ops: 0 - 0% -
Total Ops: 8 -

Node Value Delta Percent Ops Delta
node1 ---------- all ----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 5% -
Close Ops: 8 - 3% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 26% -
Getfh Ops: 22 - 7% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 18 - 6% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 10 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 92 - 31% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -
Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 4% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 9 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 294 -

```
:leveloffset: -1


= statistics node commands

:leveloffset: +1


[[IDfb4206307788f62019d2183f6898f4ce]]
```

```
= statistics node show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
System utilization metrics for each node in the cluster

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continuously displays performance data for nodes at a regular
interval. The command output displays data in the following columns:

* Node - node name.
* CPU (%) - CPU utilization.
* Total Ops - total number of operations per second.
* Latency(us) - average latency for an operation in microseconds.



== Parameters

``[-node {<nodename>|local}]`` - Node::
Selects the node for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of aggregates to display. The default setting
is 25.

== Examples
```

```
The following example displays node statistics:
```

cluster1::> statistics node show
cluster1 : 12/31/2013 16:00:04
CPU *Total Latency
Node (%) Ops (us)
----- --- ------ -------
node2 76 113 -
node1 58 10 -

```
:leveloffset: -1


= statistics oncrpc commands

:leveloffset: +1


[[IDa03ea647255c1ff371630d267b0cfbf0]]
= statistics oncrpc show-rpc-calls
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display ONC RPC Call Statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description


NOTE: This command is deprecated and will be removed in a future major
release.



The `statistics oncrpc show-rpc-calls` command displays information about
the Open Network Computing Remote Procedure Call (ONC RPC) calls performed
by the nodes of a cluster.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the RPC calls
performed by the node you specify.

``[-protocol {TCP|UDP|RDMA}]`` - Transport Protocol [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the RPC calls
performed using the network protocol you specify.

``[-badproc <Counter with Delta>]`` - Bad Procedure Calls [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the RPC calls that
have the number of bad procedure calls you specify. Bad procedure calls
are RPC requests that contain invalid procedure numbers and cannot be
completed.

``[-badlen <Counter with Delta>]`` - Bad Length Calls [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the RPC calls that
have the number of bad length calls you specify.

``[-badhdr <Counter with Delta>]`` - Bad Header Calls [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the RPC calls that
have the number of bad header calls you specify.

``[-badcalls <Counter with Delta>]`` - Bad Calls [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the RPC calls that
have the number of bad calls you specify.

``[-badprogcalls <Counter with Delta>]`` - Bad Program Calls [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the RPC calls that
have the number of bad program calls you specify.

``[-calls <Counter64 with Delta>]`` - Total Calls [.advaced-
priv]#(privilege: advanced)#::
```

```
Use this parameter to display information only about the RPC calls that
have the total number of bad calls you specify.

== Examples
```

cluster1::> statistics oncrpc show-rpc-calls

Node Value Delta
node1 ----------tcp----------
Bad Proc: 0 -
Bad Len: 0 -
Bad Hdr: 0 -
Bad Calls: 0 -
Bad Prog Calls: 0 -
Total Calls: 0 -

Node Value Delta
node1 ----------udp----------
Bad Proc: 0 -
Bad Len: 0 -
Bad Hdr: 0 -
Bad Calls: 0 -
Bad Prog Calls: 0 -
Total Calls: 0 -

Node Value Delta
node2 ----------tcp----------
Bad Proc: 0 -
Bad Len: 0 -
Bad Hdr: 0 -
Bad Calls: 0 -
Bad Prog Calls: 0 -
Total Calls: 0 -

Node Value Delta
node2 ----------udp----------
Bad Proc: 0 -
Bad Len: 0 -
Bad Hdr: 0 -
Bad Calls: 0 -
Bad Prog Calls: 0 -
Total Calls: 0 -
4 entries were displayed.

```
:leveloffset: -1


= statistics port commands

:leveloffset: +1
```

```
[[ID6860b89ba624f16e4d917016672bb785]]
= statistics port fcp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
FCP port interface throughput and latency metrics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continuously displays performance data for FCP ports at a
regular interval. The command output displays data in the following
columns:

* Port - port name.
* Read Ops - read operations per second.
* Write Ops - write operations per second.
* Other Ops - other operations per second.



== Parameters

``[-port <text>]`` - Port::
Selects the port for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of ports to display. The default setting is
```

```
25.

== Examples

The following example displays port statistics:
```

cluster1::> statistics port fcp show
cluster1 : 12/31/2013 16:00:04

```
      *Total Read Write
Port     Ops  Ops   Ops
----- ------ ---- -----
port1      2    2     0
port2      3    0     3
```

```
:leveloffset: -1


= statistics preset commands

:leveloffset: +1


[[ID4154b30df431346f08d58f1f7cccc327]]
= statistics preset delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing Performance Preset

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Deletes a performance preset configuration and all of its associated
details.

== Parameters

``-preset <text>`` - Preset Name [.advaced-priv]#(privilege: advanced)#::
Specifies the name of the performance presets that you want to delete.

== Examples
```

cluster1::*> statistics preset show
Preset Name: asup-event
Preset UUID: 55c03699-01db-11e2-8e3e-123478563412
Comment: The event-based AutoSupport Data ONTAP Performance Archive
preset configuration. This preset configuration is used
whenever an event-based AutoSupport is triggered.
Privilege: diagnostic
Read-Only: true
Archive Enabled: false
Generation ID: 0
Preset Name: asup-hourly
Preset UUID: 56178a2a-01db-11e2-8e3e-123478563412
Comment: The hourly AutoSupport Data ONTAP Performance Archive
preset configuration. This preset configuration is used by

```
the hourly AutoSupport collection events.
Privilege: diagnostic
Read-Only: true
Archive Enabled: true
Generation ID: 0
Preset Name: default
Preset UUID: 55ac6297-01db-11e2-8e3e-123478563412
Comment: The default Data ONTAP Performance Archive preset
configuration. This preset configuration includes
essential counters to assist in general troubleshooting of
system performance.
Privilege: diagnostic
Read-Only: true
Archive Enabled: true
Generation ID: 0
Preset Name: diagnostic
Preset UUID: 561db291-01db-11e2-8e3e-123478563412
Comment: The diagnostic Data ONTAP Performance Archive preset
configuration. This preset configuration includes more
counters at faster sample periods than the default
configuration to assist in troubleshooting abnormal
system performance.
Privilege: diagnostic
Read-Only: true
Archive Enabled: false
Generation ID: 0
Preset Name: foo
Preset UUID: 7a04f19d-02a7-11e2-8e40-123478563412
Comment: Test preset
Privilege: diagnostic
Read-Only: false
Archive Enabled: false
Generation ID: 0

5 entries were displayed.

cluster1::*> statistics preset delete -preset foo

cluster1::*> statistics preset show
Preset Name: asup-event
Preset UUID: 55c03699-01db-11e2-8e3e-123478563412
Comment: The event-based AutoSupport Data ONTAP Performance Archive
preset configuration. This preset configuration is used
whenever an event-based AutoSupport is triggered.
Privilege: diagnostic
Read-Only: true
Archive Enabled: false
Generation ID: 0
Preset Name: asup-hourly
Preset UUID: 56178a2a-01db-11e2-8e3e-123478563412
Comment: The hourly AutoSupport Data ONTAP Performance Archive
preset configuration. This preset configuration is used by
the hourly AutoSupport collection events.
Privilege: diagnostic
```

Read-Only: true
Archive Enabled: true
Generation ID: 0
Preset Name: default
Preset UUID: 55ac6297-01db-11e2-8e3e-123478563412
Comment: The default Data ONTAP Performance Archive preset
configuration. This preset configuration includes
essential counters to assist in general troubleshooting of
system performance.
Privilege: diagnostic
Read-Only: true
Archive Enabled: true
Generation ID: 0
Preset Name: diagnostic
Preset UUID: 561db291-01db-11e2-8e3e-123478563412
Comment: The diagnostic Data ONTAP Performance Archive preset
configuration. This preset configuration includes more
counters at faster sample periods than the default
configuration to assist in troubleshooting abnormal
system performance.
Privilege: diagnostic
Read-Only: true
Archive Enabled: false
Generation ID: 0

4 entries were displayed.

```
[[ID5b4ddcd5d018d7a1549fbce00d50a087]]
= statistics preset modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an existing Performance Preset

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Modifies an existing performance preset configuration. The command
modifies the global properties of a preset, but does not modify the
details of the preset, such as specific objects and counters sampled.

== Parameters

``-preset <text>`` - Preset Name [.advaced-priv]#(privilege: advanced)#::
Name of the performance preset to be modified.

``[-new-name <text>]`` - New Preset Name [.advaced-priv]#(privilege:
advanced)#::
Set preset name to the given new name.

``[-comment <text>]`` - Preset Description [.advaced-priv]#(privilege:
advanced)#::
Set comment to the given value.

``[-privilege <PrivilegeLevel>]`` - Preset Privilege Level [.advaced-
priv]#(privilege: advanced)#::
Set privilege level at which this preset can be viewed or modified to the
given value. Possible values: admin, advanced, diagnostic.

== Examples
```

cluster1::*> statistics preset show
Preset Name: delta
Preset UUID: 7a04f19d-02a7-11e2-8e40-123478563412
Comment: custom preset description
Privilege: diagnostic
Read-Only: false
Archive Enabled: false

Generation ID: 0

1 entry was displayed.

cluster1::*> statistics preset modify -preset delta -comment "new comment"

cluster1::*> statistics preset show
Preset Name: delta
Preset UUID: 7a04f19d-02a7-11e2-8e40-123478563412
Comment: new comment
Privilege: diagnostic
Read-Only: false
Archive Enabled: false
Generation ID: 0

1 entry was displayed.

```
[[ID45e3d54cc90aebcd0defc3d7d636ee80]]
= statistics preset show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about Performance Presets

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Displays information about performance preset configurations.

== Parameters

``[-fields <fieldname>,...]``::
Selects which performance preset attributes to display.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Shows details of all attributes of performance preset configuration.

``[-preset <text>]`` - Preset Name [.advaced-priv]#(privilege:
advanced)#::
Selects the performance presets that match the specified preset name.

``[-comment <text>]`` - Preset Description [.advaced-priv]#(privilege:
advanced)#::
Selects the performance presets that match the specified comment.
```

```
``[-privilege <PrivilegeLevel>]`` - Preset Privilege Level [.advaced-
priv]#(privilege: advanced)#::
Selects the performance presets that are available with the specified
privilege.

``[-store <text>]`` - Name of Store Where Data is Saved [.advaced-
priv]#(privilege: advanced)#::
Selects the store where data is saved.

``[-is-read-only {true|false}]`` - Is Preset Read-Only? [.advaced-
priv]#(privilege: advanced)#::
Selects the performance presets that are read-only (true) or are not read-
only (false). Read-only presets cannot be modified.

== Examples
```

cluster1::*> statistics preset show
Preset Name: aggregate_overview
Preset UUID: 4ec14b86-34a0-11ef-8db4-005056bb1920
Comment: This preset configuration is used by statistics aggregate
show command. Provides overview of aggregate object.
Privilege: admin
Read-Only: true
Archive Enabled: false
Store Name: -
Generation ID: 0
Preset Name: disk_overview
Preset UUID: 4e7471b3-34a0-11ef-8db4-005056bb1920
Comment: This preset configuration is used by statistics disk show
command. Provides overview of disk object.
Privilege: advanced
Read-Only: true
Archive Enabled: false
Store Name: -
Generation ID: 0
Preset Name: fcp_port_overview
Preset UUID: 4e77ed6e-34a0-11ef-8db4-005056bb1920
Comment: This preset configuration is used by statistics port fcp
show command. Provides overview of fcp port object.
Privilege: admin
Read-Only: true
Archive Enabled: false
Store Name: -
Generation ID: 0
Preset Name: flash_pool_overview
Preset UUID: 4ecd2576-34a0-11ef-8db4-005056bb1920
Comment: This preset configuration is used by statistics cache
flash-pool show command. Provides overview of flash pool
object.
Privilege: admin

Read-Only: true
Archive Enabled: false
Store Name: -
Generation ID: 0

```
[[IDa1a0e9bf03ccbe07a30eba4377327b89]]
= statistics preset detail show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about Performance Preset Details

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Displays the specific details of each preset, including the objects
sampled, the counter sample periods, and the counters sampled.

== Parameters

``[-fields <fieldname>,...]``::
Selects which performance preset detail attributes to display.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Displays all of the performance preset detail attributes.

``[-preset <text>]`` - Preset Name [.advaced-priv]#(privilege:
advanced)#::
Selects the performance preset details that match the specified preset
name.

``[-object <text>]`` - Performance Object [.advaced-priv]#(privilege:
advanced)#::
Selects the performance preset details that match the specified object
name.

``[-sample-period <sample_period>]`` - Archive Sample period [.advaced-
priv]#(privilege: advanced)#::
Selects the performance preset details that are collected at the specified
sample period.

``[-counter-set <text>,...]`` - Performance Counter Name Set [.advaced-
```

```
priv]#(privilege: advanced)#::
Selects the performance preset details that match the specified counters
in the counter set. Use "|" to separate multiple counters.

``[-instance-filters <text>,...]`` - Performance Instance Filters
[.advaced-priv]#(privilege: advanced)#::
Selects the performance preset details that match the specified instance
filters. Use "|" to separate multiple instance filters. This field is
reserved for future use.

== Examples
```

cluster1::*> statistics preset detail show
Sample Counter Instance
Preset Name Object Period Set Filters
---------- ---------- ------ ---------- ----------
asup-event aggregate 1w instance_name, -
node_name,
process_name,
parent_host,
total_
transfers,
user_reads,
user_writes,
cp_reads,
user_read_
blocks,
user_write_
blocks,
cp_read_
blocks,
wv_fsid,
wv_vol_type,
wv_fsinfo_fs_
version,
wv_volinfo_fs_
options,
...

```
:leveloffset: -1


= statistics qtree commands

:leveloffset: +1


[[ID408c0d778f7fd70d85efeaeb64cdd9c9]]
= statistics qtree show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Qtree I/O operation rates

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command continuously displays performance data for qtrees at a
regular interval. The command output displays data in the following
columns:

* Qtree - Qtree name.
* Vserver - Vserver name.
* Volume - Volume name.
* NFS Ops - NFS operations per second.
* CIFS Ops - CIFS operations per second.
* Internal Ops - Internal operations per second.
* Total Ops - Total number of operations per second.



== Parameters

``[-qtree <text>]`` - Qtree::
Selects the qtree for which you want to display performance data. The
qtree name has to be given in the format "vol_name/qtree_name"

``[-vserver <vserver name>]`` - Vserver::
Selects the Vserver for which you want to display performance data.

``[-volume <text>]`` - Volume::
Selects the volume for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.
```

```
``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of qtrees to display. The default setting is
25.

== Examples

The following example displays qtree statistics:
```

C1_sti20-vsim-ucs429g_1520278254::> statistics qtree show
C1_sti20-vsim-ucs429g_1520278254 : 3/14/2018 15:40:46
NFS CIFS Internal *Total
Qtree Vserver Volume ops ops ops ops
---------- ------- ------- --- ---- -------- ------
flexvol/qt1 vs0 flexvol 7 0 0 7
dp_vol/qt9 vs0 dp_vol 7 0 0 7
dp_vol/qt8 vs0 dp_vol 7 0 0 7

```
:leveloffset: -1


= statistics samples commands

:leveloffset: +1


[[IDe55179b81b34413edba09e4a9acfff00]]
= statistics samples delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete statistics samples

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command deletes samples that you created using the
xref:{relative_path}statistics-start.html[statistics start] command.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Selects the Vserver for which you want to delete the sample. The default
Vserver is admin Vserver.

``-sample-id <text>`` - Sample Identifier [.advaced-priv]#(privilege:
advanced)#::
Specifies the sample that you want to delete. This is a required
parameter.

== Examples

The following example deletes the sample "smpl_1":
```

cluster1::*> statistics samples delete -sample-id smpl_1

```
== Related Links
```

```
* xref:{relative_path}statistics-start.html[statistics start]




[[ID77ec8039764b276c93d94c6ce667601b]]
= statistics samples show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display statistics samples

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command displays information about the samples that you created using
the xref:{relative_path}statistics-start.html[statistics start] command.

== 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.

| ``[-describe  ]`` [.advaced-priv]#(privilege: advanced)# }::
Displays detailed information about each sample.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Selects the samples that match this parameter value. If you omit this
parameter, the command displays details for all samples.

``[-sample-id <text>]`` - Sample Identifier [.advaced-priv]#(privilege:
advanced)#::
Selects the samples that match this parameter value. If you do not specify
this parameter, the command will display information about all the samples
in the cluster.

== Examples
```

```
The following example displays information for sample "smpl_1":
```

cluster1::*> statistics samples show -sample-id smpl_1
Vserver Sample ID Start Time Stop Time Status
--- --- --- --- ---
cluster-d1 smpl_1 09/13 18:06:46 - Ready

```
The following example displays detailed information for sample "smpl_1":
```

cluster1::*> statistics samples show -sample-id smpl_1 -describe
Vserver: vs1
Sample ID: smpl_1
Object: processor
Instance: -
Counter: -
Start Time: 09/13 18:06:46
Stop Time: -
Status: Ready - -
Privilege: admin

```
== Related Links

* xref:{relative_path}statistics-start.html[statistics start]




:leveloffset: -1


= statistics settings commands

:leveloffset: +1


[[IDa49e5cf74dd0dfd714e65530a95f3ec4]]
= statistics settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify settings for the statistics commands

*Availability:* This command is available to _cluster_ administrators at
```

```
the _advanced_ privilege level.

== Description

This command modifies the settings for all of the `statistics` commands.

== Parameters

``[-display-rates {true|false}]`` - Display Rates [.advaced-
priv]#(privilege: advanced)#::
Specifies whether the `statistics` commands display rate counters in
rates/second. The default is true.

``[-client-stats {enabled|disabled}]`` - Collect Per-Client Statistics
[.advaced-priv]#(privilege: advanced)#::
Specifies whether `statistics` commands display per-client information.
The default is disabled.
+
NOTE: If you enable this setting, you might significantly impact system
performance.
+


``[-counter-display-units {B|KB|MB|GB}]`` - Counter Display Units
[.advaced-priv]#(privilege: advanced)#::
Specifies display units for the counters. The default setting is MB.

``[-display-count-exponent <integer>]`` - Display Count Exponent
[.advaced-priv]#(privilege: advanced)#::
Specifies display exponent value for the counters representing counts. The
default setting is 3 (thousand).

== Examples

The following example sets the value of the `-display-rates` parameter to
false :
```

cluster1::*> statistics settings modify -display-rates false

```
[[ID5de24e792dfc243c0b2cce439e2ce324]]
= statistics settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display settings for the statistics commands

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command displays the current settings for all of the `statistics`
commands.

== Examples

The following example displays the current settings for all `statistics`
commands:
```

cluster1::*> statistics settings show
Display rate Counters in rate/sec: true

Counter Display: full
Counter Display Units: MB
Display Count Exponent: 3

```
:leveloffset: -1


= statistics system commands

:leveloffset: +1


[[ID9938cbcd35d1b84475b2459253c4d0af]]
= statistics system show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
System utilization metrics for the cluster

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continuously displays performance data for cluster at a
regular interval. The command output displays data in the following
columns:

* System - cluster name.
* CPU (%) - CPU utilization.
* Total Ops - total number of operations per second.
* Latency(ms) - average latency for an operation in milliseconds.



== Parameters

``[-system <text>]`` - System::
Selects the cluster for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of systems to display. The default setting is
25.

== Examples

The following example displays system statistics:
```

cluster1::> statistics system show
cluster1 : 12/31/2013 16:00:04

```
        CPU *Total Latency
 System (%)    Ops    (ms)
------- --- ------ -------
Cluster  76    113       -
```

```
:leveloffset: -1


= statistics top commands

:leveloffset: +1


[[IDbb60c64b63711af175159400e8d489af]]
= statistics top client show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Most active NFS and CIFS clients

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continuously displays performance data for the top NFS and
CIFS clients at a regular interval. The command output displays data in
the following columns:

* Client - client name.
* Vserver - vserver name.
* Node - node name.
* Total IOPs - total number of operations per second.
* Total (Bps) - total throughput in bytes per second.



== Parameters

``[-node {<nodename>|local}]`` - Node::
Selects the node for which you want to display performance data.
```

```
``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies maximum number of top clients to display. The default setting is
10.

== Examples

The following example displays top client statistics:
```

cluster1::> statistics top client show
cluster-1 : 12/31/1969 16:00:04
*Total Total
Client Vserver Node Ops (Bps)
--- --- --- --- ---
172.17.236.53:938 vserver01 cluster-node2 9 80
172.17.236.160:898 vserver02 cluster-node1 6 50

```
[[ID3f7965887dca6c7a9b2679f4f3712d02]]
= statistics top file show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Most actively accessed files

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continuously displays performance data for top files at a
regular interval. The command output displays data in the following
```

```
columns:

* Total IOPS - total number of input/output operations per second.
* Node - node name.
* Vserver - vserver name.
* Aggregate - aggregate name.
* Volume - volume name.
* File - file name.



== Parameters

``[-node {<nodename>|local}]`` - Node::
Selects the node for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 30 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies maximum number of top files to display. The default setting is
10.

== Examples

The following example displays top files statistics:
```

cluster1::> statistics top file show
cluster-1 : 12/31/1969 16:00:04
*Estimated
Total
IOPS Node Vserver Aggregate Volume File
---------- ------------- --------- --------- ------ ---------------------
9 cluster-node2 vserver01 aggr1 vol01 /vol/vol01/clus/cache
6 cluster-node1 vserver02 aggr2 vol02 /vol/vol02

```
:leveloffset: -1
```

```
= statistics volume commands

:leveloffset: +1


[[ID2af496381a6ca671cbe50eb30b2abf3f]]
= statistics volume show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Volume throughput and latency metrics

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command continuously displays performance data for volumes at a
regular interval. The command output displays data in the following
columns:

* Volume - volume name.
* Vserver - vserver name.
* Aggregate - aggregate name.
* Total Ops - total number of operations per second.
* Read Ops - read operations per second.
* Write Ops - write operations per second.
* Other Ops - other operations per second.
* Read (Bps) - read throughput in bytes per second.
* Write (Bps) - write throughput in bytes per second.
* Latency(us) - average latency for an operation in microseconds.



== Parameters

``[-volume <text>]`` - Volume::
Selects the volume for which you want to display performance data.

``[-vserver <vserver name>]`` - Vserver::
Selects the vserver for which you want to display performance data.
```

```
``[-aggregate <text>]`` - Aggregate::
Selects the aggregate for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of volumes to display. The default setting is
25.

== Examples

The following example displays volume statistics:
```

cluster1::> statistics volume show
cluster1 : 12/31/2013 16:00:04

```
                          *Total Read Write Other   Read Write Latency
Volume Vserver Aggregate    Ops  Ops   Ops   Ops  (Bps) (Bps)    (us)
------ ------- --------- ------ ---- ----- ----- ------ ----- -------
  vol0       -     aggr0     58   13    15    29 310585  3014      39
  vol0       -  aggr0_n0     56    0    11    45   8192 28826      47
```

```
:leveloffset: -1


= statistics vserver commands

:leveloffset: +1


[[ID6bb5fdd16e55fa5c244e5e1a25f43c67]]
= statistics vserver show
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Vserver throughput and latency metrics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continuously displays performance data for Vservers at a
regular interval. The command output displays data in the following
columns:

* Vserver - Vserver name.
* Total Ops - total number of operations per second.
* Read Ops - read operations per second.
* Write Ops - write operations per second.
* Other Ops - other operations per second.
* Read (Bps) - read throughput in bytes per second.
* Write (Bps) - write throughput in bytes per second.
* Latency(us) - average latency for an operation in microseconds.



== Parameters

``[-vserver <vserver name>]`` - Vserver::
Selects the vserver for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances::
Specifies the maximum number of Vservers to display. The default setting
```

```
is 25.

== Examples

The following example displays Vserver statistics:
```

cluster1::> statistics vserver show
cluster1 : 12/31/2013 16:00:04

```
        *Total Read Write Other   Read Write Latency
Vserver    Ops  Ops   Ops   Ops  (Bps) (Bps)    (us)
------- ------ ---- ----- ----- ------ ----- -------
    vs1     58   13    15    29 310585  3014      39
    vs2     56    0    11    45   8192 28826      47
```

```
:leveloffset: -1


= statistics workload commands

:leveloffset: +1


[[ID46324acdcbda04cf91a125475240f1e1]]
= statistics workload show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
QoS workload throughput and latency metrics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command continuously displays performance data for workloads at a
regular interval. The command output displays data in the following
columns:

* Workload - workload name.
* Total Ops - total number of operations per second.
```

```
* Read Ops - read operations per second.
* Write Ops - write operations per second.
* Other Ops - other operations per second.
* Read (Bps) - read throughput in bytes per second.
* Write (Bps) - write throughput in bytes per second.
* Latency(us) - average latency for an operation in microseconds.



== Parameters

``[-workload <text>]`` - Workload [.advaced-priv]#(privilege: advanced)#::
Selects the workload for which you want to display performance data.

``[-sort-key <text>]`` - Column to Sort By [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays statistics sorted by
the specified column.

``-interval <integer>`` - Interval [.advaced-priv]#(privilege:
advanced)#::
Specifies, in seconds, the interval between statistics updates. The
default setting is 5 seconds.

``-iterations <integer>`` - Iterations [.advaced-priv]#(privilege:
advanced)#::
Specifies the number of iterations the command runs before terminating.
The default setting is 1. If the number is 0 (zero), the command continues
to run until you interrupt it by pressing Ctrl-C.

``-max <integer>`` - Maximum Number of Instances [.advaced-
priv]#(privilege: advanced)#::
Specifies the maximum number of workloads to display. The default setting
is 25.

== Examples

The following example displays workload statistics:
```

cluster1::> statistics workload show
cluster1 : 12/31/2013 16:00:04

```
                 *Total Read Write Other   Read Write Latency
       Workload    Ops  Ops   Ops   Ops  (Bps) (Bps)    (us)
--------------- ------ ---- ----- ----- ------ ----- -------
_USERSPACE_APPS     30    1     3     0  30765  8553       0
     _WAFL_SCAN     20    0     0     0      0     0       0
       _WAFL_CP      0    0     0     0      0     0       -
```

```
:leveloffset: -1


:leveloffset: -1


= statistics-v1 commands

:leveloffset: +1


= statistics-v1 nfs commands

:leveloffset: +1


[[IDd610bcfb6f704e1832f1ee3845059437]]
= statistics-v1 nfs show-mount
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display mount statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `statistics-v1 nfs show-mount` command displays the following
statistics about the NFS mounts on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of mount operations
* Total number of dump operations
```

```
* Total number of unmount operations
* Total number of unmountall operations
* Total number of export operations
* Total number of exportall operations
* Total number of pathconf operations
* Total number of all the above operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only for
the specified node.

``[-result {success|failure|all}]`` - Result [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-mount <Counter with Delta>]`` - Mount Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of mount operations.

``[-dump <Counter with Delta>]`` - Dump Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of dump operations.

``[-unmnt <Counter with Delta>]`` - UnMount Operations [.advaced-
```

```
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmount operations.

``[-unmntall <Counter with Delta>]`` - UnMountAll Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmountall operations.

``[-export <Counter with Delta>]`` - Export Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of export operations.

``[-exportall <Counter with Delta>]`` - ExportAll Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of exportall operations.

``[-pathconf <Counter with Delta>]`` - PathConf Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of pathconf operations.

``[-total <Counter64 with Delta>]`` - Total Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total operations.

== Examples

The following example displays statistics about the NFS mounts for a node
named node1:
```

cluster1::*> statistics-v1 nfs show-mount -node node1

Node Value Delta
node1 ----------success-------
Null Ops: 2 0/s:16s
Mount Ops: 1 0/s:16s
Dump Ops: 0 -
Unmount Ops: 1 0/s:16s
Unmount All Ops: 0 -
Export Ops: 0 -
ExportAll Ops 0 -
PathConf Ops: 0 -
Total Ops: 4 0/s:16s

Node Value Delta
node1 ----------failure-------
Null Ops: 0 -
Mount Ops: 0 -
Dump Ops: 0 -
Unmount Ops: 0 -
Unmount All Ops: 0 -
Export Ops: 0 -
ExportAll Ops 0 -
PathConf Ops: 0 -
Total Ops: 0 -

```
[[ID6d9399b96d2f2e166036745c12c1d056]]
= statistics-v1 nfs show-nlm
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Display NLM statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `statistics-v1 nfs show-nlm` command displays the following statistics
about the Network Lock Manager (NLM) on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of test operations
* Total number of lock operations
* Total number of cancel operations
* Total number of unlock operations
* Total number of granted operations
* Total number of share operations
* Total number of unshare operations
* Total number of nmlock operations
* Total number of freeall operations
* Total number of all the above operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

NOTE: This command requires an effective cluster version earlier than
ONTAP 9.0. Data for nodes running ONTAP 9.0 or later is not collected, and
```

```
will not be displayed. Use the xref:{relative_path}statistics-
show.html[statistics show]`-object```_nlm_`` command instead.



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only for
the specified node.

``[-result {success|failure|all}]`` - Result [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-test <Counter with Delta>]`` - Test Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of test operations.

``[-lock <Counter with Delta>]`` - Lock Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lock operations.

``[-cancel <Counter with Delta>]`` - Cancel Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of cancel operations.

``[-unlock <Counter with Delta>]`` - Unlock Operations [.advaced-
```

```
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unlock operations.

``[-granted <Counter with Delta>]`` - Granted Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of granted operations.

``[-share <Counter with Delta>]`` - Share Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of share operations.

``[-unshare <Counter with Delta>]`` - Unshare Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unshare operations.

``[-nmlock <Counter with Delta>]`` - NmLock Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of nmlock operations.

``[-freeall <Counter with Delta>]`` - FreeAll Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of freeall operations.

``[-total <Counter64 with Delta>]`` - Total Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total operations.

== Examples

The following example displays statistics about the NLM for a node named
node1:
```

cluster1::*> statistics-v1 nfs show-nlm -node node1

Node Value Delta
node1 --------success-------
Null: 0 -
Test: 0 -
Lock: 2 0/s:23s
Cancel: 0 -

Unlock: 1 0/s:23s
Granted: 0 -
Share: 0 -
Unshare: 0 -
NmLock: 0 -
FreeAll: 0 -
Total: 3 0/s:23s

Node Value Delta
node1 --------failure-------
Null: 0 -
Test: 0 -
Lock: 0 -
Cancel: 0 -
Unlock: 0 -
Granted: 0 -
Share: 0 -
Unshare: 0 -
NmLock: 0 -
FreeAll: 0 -
Total: 0 -

```
== Related Links

* xref:{relative_path}statistics-show.html[statistics show]




[[ID9d57adcb566a2c6f5c8187de98afec76]]
= statistics-v1 nfs show-statusmon
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status monitor statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `statistics-v1 nfs show-statusmon` command displays the following
statistics about the Status Monitor on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
```

```
* Total number of stat operations
* Total number of monitor operations
* Total number of unmonitor operations
* Total number of unmonitor all operations
* Total number of simucrash operations
* Total number of notify operations
* Total number of all the above operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only for
the specified node.

``[-result {success|failure|all}]`` - Result [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-stat <Counter with Delta>]`` - Stat Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of stat operations.

``[-monitor <Counter with Delta>]`` - Monitor Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of monitor operations.
```

```
``[-unmonitor <Counter with Delta>]`` - Unmonitor Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmonitor operations.

``[-unmonall <Counter with Delta>]`` - Unmonitor All Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of unmonitor all
operations.

``[-simucrash <Counter with Delta>]`` - SimuCrash Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of simucrash operations.

``[-notify <Counter with Delta>]`` - Notify Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of notify operations.

``[-total <Counter64 with Delta>]`` - Total Operations [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total operations.

== Examples

The following example displays statistics about the status monitor for a
node named node1:
```

cluster1::*> statistics-v1 nfs show-statusmon -node node1

Node Value Delta
node1 --------success-------
Null Ops: 0 -
Stat Ops: 0 -
Monitor Ops: 0 -
Unmonitor Ops: 0 -
Unmon All Ops: 0 -
SimuCrash Ops: 0 -
Notify Ops: 0 -
Total Ops: 0 -

Node Value Delta
node1 --------failure-------
Null Ops: 0 -
Stat Ops: 0 -

Monitor Ops: 0 -
Unmonitor Ops: 0 -
Unmon All Ops: 0 -
SimuCrash Ops: 0 -
Notify Ops: 0 -
Total Ops: 0 -

```
[[IDf361f2eb169ca93b1ce825197f3c1129]]
= statistics-v1 nfs show-v3
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NFSv3 statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `statistics-v1 nfs show-v3` command displays the following statistics
about the NFSv3 operations on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of getattr operations
* Total number of setattr operations
* Total number of lookup operations
* Total number of access operations
* Total number of readsymlink operations
* Total number of read operations
* Total number of write operations
* Total number of create operations
* Total number of mkdir operations
* Total number of symlink operations
* Total number of mknod operations
* Total number of remove operations
* Total number of rmdir operations
* Total number of rename operations
* Total number of link operations
* Total number of readdir operations
* Total number of readdirplus operations
* Total number of fsstat operations
* Total number of fsinfo operations
* Total number of pathconf operations
```

```
* Total number of commit operations
* Total number of nfsv3 operations
* Percent of null operations
* Percent of getattr operations
* Percent of setattr operations
* Percent of lookup operations
* Percent of access operations
* Percent of readsymlink operations
* Percent of read operations
* Percent of write operations
* Percent of create operations
* Percent of mkdir operations
* Percent of symlink operations
* Percent of mknod operations
* Percent of remove operations
* Percent of rmdir operations
* Percent of rename operations
* Percent of link operations
* Percent of readdir operations
* Percent of readdirplus operations
* Percent of fsstat operations
* Percent of fsinfo operations
* Percent of pathconf operations
* Percent of commit operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays NFSv3 statistics only
for the specified node.

``[-result {success|failure|all}]`` - Result::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).
```

```
``[-null <Counter with Delta>]`` - Null Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-gattr <Counter with Delta>]`` - GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getattr operations.

``[-sattr <Counter with Delta>]`` - SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setattr operations.

``[-lookup <Counter with Delta>]`` - LookUp Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookup operations.

``[-access <Counter with Delta>]`` - Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of access operations.

``[-rsym <Counter with Delta>]`` - ReadSymlink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readsymlink
operations.

``[-read <Counter with Delta>]`` - Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of read operations.

``[-write <Counter with Delta>]`` - Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of write operations.

``[-create <Counter with Delta>]`` - Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of create operations.

``[-mkdir <Counter with Delta>]`` - MkDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of mkdir operations.

``[-symln <Counter with Delta>]`` - SymLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of symlink operations.

``[-mknod <Counter with Delta>]`` - MkNod Operations::
If you specify this parameter, the command displays statistics only about
```

```
the node or nodes that have the specified number of mknod operations.

``[-remove <Counter with Delta>]`` - Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of remove operations.

``[-rmdir <Counter with Delta>]`` - RmDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rmdir operations.

``[-rename <Counter with Delta>]`` - Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rename operations.

``[-link <Counter with Delta>]`` - Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of link operations.

``[-rdir <Counter with Delta>]`` - ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readdir operations.

``[-rdirp <Counter with Delta>]`` - ReadDirPlus Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readdirplus
operations.

``[-fsstat <Counter with Delta>]`` - FsStat Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of fsstat operations.

``[-fsinfo <Counter with Delta>]`` - FsInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of fsinfo operations.

``[-pconf <Counter with Delta>]`` - PathConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of pathconf operations.

``[-commit <Counter with Delta>]`` - Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of commit operations.

``[-total <Counter64 with Delta>]`` - Total Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total NFSv3
operations.
```

```
``[-null-pct <Counter with Delta>]`` - Percent Null Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of null operations.

``[-gattr-pct <Counter with Delta>]`` - Percent GetAttr Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getattr
operations.

``[-sattr-pct <Counter with Delta>]`` - Percent SetAttr Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setattr
operations.

``[-lookup-pct <Counter with Delta>]`` - Percent LookUp Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookup operations.

``[-access-pct <Counter with Delta>]`` - Percent Access Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of access operations.

``[-rsym-pct <Counter with Delta>]`` - Percent ReadSymlink Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readsymlink
operations.

``[-read-pct <Counter with Delta>]`` - Percent Read Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of read operations.

``[-write-pct <Counter with Delta>]`` - Percent Write Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of write operations.

``[-create-pct <Counter with Delta>]`` - Percent Create Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of create operations.

``[-mkdir-pct <Counter with Delta>]`` - Percent MkDir Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of mkdir operations.

``[-symln-pct <Counter with Delta>]`` - Percent SymLink Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of symlink
```

```
operations.

``[-mknod-pct <Counter with Delta>]`` - Percent MkNod Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of mknod operations.

``[-remove-pct <Counter with Delta>]`` - Percent Remove Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of remove operations.

``[-rmdir-pct <Counter with Delta>]`` - Percent RmDir Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rmdir operations.

``[-rename-pct <Counter with Delta>]`` - Percent Rename Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rename operations.

``[-link-pct <Counter with Delta>]`` - Percent Link Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of link operations.

``[-rdir-pct <Counter with Delta>]`` - Percent ReadDir Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdir
operations.

``[-rdirp-pct <Counter with Delta>]`` - Percent ReadDirPlus Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdirplus
operations.

``[-fsstat-pct <Counter with Delta>]`` - Percent FsStat Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of fsstat operations.

``[-fsinfo-pct <Counter with Delta>]`` - Percent FsInfo Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of fsinfo operations.

``[-pconf-pct <Counter with Delta>]`` - Percent PathConf Ops::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of pathconf
operations.

``[-commit-pct <Counter with Delta>]`` - Percent Commit Ops::
If you specify this parameter, the command displays statistics only about
```

```
the node or nodes that have the specified percentage of commit operations.

== Examples

The following example displays statistics about the NFSv3 operations for a
node named node1:
```

cluster1::> statistics-v1 nfs show-v3 -node node1

Node Value Delta Percent Ops Delta
node1 ----------success----------
Null Ops: 4 - 7% -
GetAttr Ops: 10 - 19% -
SetAttr Ops: 2 - 4% -
Lookup Ops: 2 - 4% -
Access Ops: 14 - 26% -
ReadSymlink Ops: 0 - 0% -
Read Ops: 0 - 0% -
Write Ops: 0 - 0% -
Create Ops: 2 - 4% -
MkDir Ops: 1 - 2% -
Symlink Ops: 0 - 0% -
MkNod Ops: 0 - 0% -
Remove Ops: 1 - 2% -
RmDir Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Link Ops: 0 - 0% -
ReadDir Ops: 2 - 4% -
ReadDirPlus Ops: 10 - 19% -
FsStat Ops: 1 - 2% -
FsInfo Ops: 5 - 9% -
PathConf Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Total Ops: 54 -

Node Value Delta Percent Ops Delta
node1 ----------failure----------
Null Ops: 0 - 0% -
GetAttr Ops: 0 - 0% -
SetAttr Ops: 0 - 0% -
Lookup Ops: 2 - 100% -
Access Ops: 0 - 0% -
ReadSymlink Ops: 0 - 0% -
Read Ops: 0 - 0% -
Write Ops: 0 - 0% -
Create Ops: 0 - 0% -
MkDir Ops: 0 - 0% -
Symlink Ops: 0 - 0% -
MkNod Ops: 0 - 0% -
Remove Ops: 0 - 0% -
RmDir Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Link Ops: 0 - 0% -

ReadDir Ops: 0 - 0% -
ReadDirPlus Ops: 0 - 0% -
FsStat Ops: 0 - 0% -
FsInfo Ops: 0 - 0% -
PathConf Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Total Ops: 2 -

Node Value Delta Percent Ops Delta
node1 ---------- all ----------
Null Ops: 4 - 7% -
GetAttr Ops: 10 - 18% -
SetAttr Ops: 2 - 4% -
Lookup Ops: 4 - 7% -
Access Ops: 14 - 25% -
ReadSymlink Ops: 0 - 0% -
Read Ops: 0 - 0% -
Write Ops: 0 - 0% -
Create Ops: 2 - 4% -
MkDir Ops: 1 - 2% -
Symlink Ops: 0 - 0% -
MkNod Ops: 0 - 0% -
Remove Ops: 1 - 2% -
RmDir Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Link Ops: 0 - 0% -
ReadDir Ops: 2 - 4% -
ReadDirPlus Ops: 10 - 18% -
FsStat Ops: 1 - 2% -
FsInfo Ops: 5 - 9% -
PathConf Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Total Ops: 56 -

```
[[ID8eac124a5eaee88e39c2fb5d7e097cdf]]
= statistics-v1 nfs show-v4
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NFSv4 statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `statistics-v1 nfs show-v4` command displays the following statistics
about the NFSv4 operations on each node in the cluster:
```

```
* Result of the operations (success or failure)
* Total number of null operations
* Total number of compound operations
* Total number of access operations
* Total number of close operations
* Total number of commit operations
* Total number of create operations
* Total number of delegpurge operations
* Total number of delegret operations
* Total number of getattr operations
* Total number of getfh operations
* Total number of link operations
* Total number of lock operations
* Total number of lockt operations
* Total number of locku operations
* Total number of lookup operations
* Total number of lookupp operations
* Total number of nverify operations
* Total number of open operations
* Total number of openattr operations
* Total number of openconf operations
* Total number of opendowng operations
* Total number of putfh operations
* Total number of putpubfh operations
* Total number of putrootfh operations
* Total number of read operations
* Total number of readdir operations
* Total number of readlink operations
* Total number of remove operations
* Total number of rename operations
* Total number of renew operations
* Total number of restorefh operations
* Total number of savefh operations
* Total number of secinfo operations
* Total number of setattr operations
* Total number of setcliid operations
* Total number of setcliidconf operations
* Total number of verify operations
* Total number of write operations
* Total number of rellockown operations
* Total number of total operations
* Percent of null operations
* Percent of compound operations
* Percent of access operations
* Percent of close operations
```

```
* Percent of commit operations
* Percent of create operations
* Percent of delegpurge operations
* Percent of delegret operations
* Percent of getattr operations
* Percent of getfh operations
* Percent of link operations
* Percent of lock operations
* Percent of lockt operations
* Percent of locku operations
* Percent of lookup operations
* Percent of lookupp operations
* Percent of nverify operations
* Percent of open operations
* Percent of openattr operations
* Percent of openconf operations
* Percent of opendowng operations
* Percent of putfh operations
* Percent of putpubfh operations
* Percent of putrootfh operations
* Percent of read operations
* Percent of readdir operations
* Percent of readlink operations
* Percent of remove operations
* Percent of rename operations
* Percent of renew operations
* Percent of restorefh operations
* Percent of savefh operations
* Percent of secinfo operations
* Percent of setattr operations
* Percent of setcliid operations
* Percent of setCliidconf operations
* Percent of verify operations
* Percent of write operations
* Percent of rellockown operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays NFSv4 statistics only
for the specified node.

``[-result {success|failure|all}]`` - Result::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-cmpnd <Counter with Delta>]`` - Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of compound operations.

``[-access <Counter with Delta>]`` - Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of access operations.

``[-close <Counter with Delta>]`` - Close Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of close operations.

``[-commit <Counter with Delta>]`` - Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of commit operations.

``[-create <Counter with Delta>]`` - Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of create operations.

``[-delpur <Counter with Delta>]`` - Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegpurge operations.

``[-delrtn <Counter with Delta>]`` - Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegret operations.

``[-gattr <Counter with Delta>]`` - GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getattr operations.
```

```
``[-getfh <Counter with Delta>]`` - GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getfh operations.

``[-link <Counter with Delta>]`` - Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of link operations.

``[-lock <Counter with Delta>]`` - Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lock operations.

``[-lockt <Counter with Delta>]`` - LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lockt operations.

``[-locku <Counter with Delta>]`` - LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of locku operations.

``[-lookup <Counter with Delta>]`` - Lookup Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookup operations.

``[-lookpp <Counter with Delta>]`` - LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookupp operations.

``[-nverfy <Counter with Delta>]`` - Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of nverify operations.

``[-open <Counter with Delta>]`` - Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of open operations.

``[-opattr <Counter with Delta>]`` - OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openattr operations.

``[-opconf <Counter with Delta>]`` - OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openconf operations.

``[-opndg <Counter with Delta>]`` - OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
```

```
the node or nodes that have the specified number of opendowng operations.

``[-putfh <Counter with Delta>]`` - PutFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putfh operations.

``[-putpfh <Counter with Delta>]`` - PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putpubfh operations.

``[-putrfh <Counter with Delta>]`` - PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putrootfh operations.

``[-read <Counter with Delta>]`` - Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of read operations.

``[-readdr <Counter with Delta>]`` - ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readdir operations.

``[-rlink <Counter with Delta>]`` - ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readlink operations.

``[-remove <Counter with Delta>]`` - Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of remove operations.

``[-rename <Counter with Delta>]`` - Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rename operations.

``[-renew <Counter with Delta>]`` - Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of renew operations.

``[-restfh <Counter with Delta>]`` - RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of restorefh operations.

``[-savefh <Counter with Delta>]`` - SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of savefh operations.

``[-secinf <Counter with Delta>]`` - SecInfo Operations::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of secinfo operations.

``[-sattr <Counter with Delta>]`` - SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setattr operations.

``[-sclid <Counter with Delta>]`` - SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliid operations.

``[-scidc <Counter with Delta>]`` - SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliidconf
operations.

``[-verify <Counter with Delta>]`` - Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of verify operations.

``[-write <Counter with Delta>]`` - Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of write operations.

``[-relown <Counter with Delta>]`` - RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rellockown operations.

``[-total <Counter64 with Delta>]`` - Total Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total nfsv4
operations.

``[-null-pct <Counter with Delta>]`` - Percent Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of null operations.

``[-cmpnd-pct <Counter with Delta>]`` - Percent Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of compound
operations.

``[-access-pct <Counter with Delta>]`` - Percent Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of access operations.

``[-close-pct <Counter with Delta>]`` - Percent Close Operations::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of close operations.

``[-commit-pct <Counter with Delta>]`` - Percent Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of commit operations.

``[-create-pct <Counter with Delta>]`` - Percent Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of create operations.

``[-delpur-pct <Counter with Delta>]`` - Percent Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegpurge
operations.

``[-delrtn-pct <Counter with Delta>]`` - Percent Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegret
operations.

``[-gattr-pct <Counter with Delta>]`` - Percent GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getattr
operations.

``[-getfh-pct <Counter with Delta>]`` - Percent GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getfh operations.

``[-link-pct <Counter with Delta>]`` - Percent Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of link operations.

``[-lock-pct <Counter with Delta>]`` - Percent Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lock operations.

``[-lockt-pct <Counter with Delta>]`` - Percent LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lockt operations.

``[-locku-pct <Counter with Delta>]`` - Percent LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of locku operations.

``[-lookup-pct <Counter with Delta>]`` - Percent Lookup Operations::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookup operations.

``[-lookpp-pct <Counter with Delta>]`` - Percent LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookupp
operations.

``[-nverfy-pct <Counter with Delta>]`` - Percent Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of nverify
operations.

``[-open-pct <Counter with Delta>]`` - Percent Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of open operations.

``[-opattr-pct <Counter with Delta>]`` - Percent OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openattr
operations.

``[-opconf-pct <Counter with Delta>]`` - Percent OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openconf
operations.

``[-opndg-pct <Counter with Delta>]`` - Percent OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of opendowng
operations.

``[-putfh-pct <Counter with Delta>]`` - Percent PutFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putfh operations.

``[-putpfh-pct <Counter with Delta>]`` - Percent PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putpubfh
operations.

``[-putrfh-pct <Counter with Delta>]`` - Percent PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putrootfh
operations.

``[-read-pct <Counter with Delta>]`` - Percent Read Operations::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of read operations.

``[-readdr-pct <Counter with Delta>]`` - Percent ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdir
operations.

``[-rlink-pct <Counter with Delta>]`` - Percent ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readlink
operations.

``[-remove-pct <Counter with Delta>]`` - Percent Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of remove operations.

``[-rename-pct <Counter with Delta>]`` - Percent Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rename operations.

``[-renew-pct <Counter with Delta>]`` - Percent Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of renew operations.

``[-restfh-pct <Counter with Delta>]`` - Percent RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of restorefh
operations.

``[-savefh-pct <Counter with Delta>]`` - Percent SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of savefh operations.

``[-secinf-pct <Counter with Delta>]`` - Percent SecInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of secinfo
operations.

``[-sattr-pct <Counter with Delta>]`` - Percent SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setattr
operations.

``[-sclid-pct <Counter with Delta>]`` - Percent SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliid
```

```
operations.

``[-scidc-pct <Counter with Delta>]`` - Percent SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliidconf
operations.

``[-verify-pct <Counter with Delta>]`` - Percent Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of verify operations.

``[-write-pct <Counter with Delta>]`` - Percent Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of write operations.

``[-relown-pct <Counter with Delta>]`` - Percent RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rellockown
operations.

== Examples

The following example displays statistics about the NFSv4 operations for a
node named node1:
```

cluster1::> statistics-v1 nfs show-v4 -node node1

Node Value Delta Percent Ops Delta
node1 ----------success----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 6% -
Close Ops: 8 - 3% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 27% -
Getfh Ops: 22 - 8% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 13 - 5% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 8 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -

Putfh Ops: 92 - 32% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -
Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 5% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 8 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 286 -
Node Value Delta Percent Ops Delta
node1 ----------failure----------
Null Procs: 0 - 0% -
Cmpnd Procs: 0 - -
Access Ops: 0 - 0% -
Close Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 0 - 0% -
Getfh Ops: 0 - 0% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 5 - 63% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 2 - 25% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 0 - 0% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 0 - 0% -
Read Ops: 0 - 0% -
Readdir Ops: 0 - 0% -
Readlink Ops: 0 - 0% -
Remove Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Renew Ops: 0 - 0% -
Restorefh Ops: 0 - 0% -
Savefh Ops: 0 - 0% -
Secinfo Ops: 0 - 0% -

Setattr Ops: 1 - 13% -
Setclid Ops: 0 - 0% -
Setclidconf Ops: 0 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 0 - 0% -
Rlockown Ops: 0 - 0% -
Total Ops: 8 -
Node Value Delta Percent Ops Delta
node1 ---------- all ----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 5% -
Close Ops: 8 - 3% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 26% -
Getfh Ops: 22 - 7% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 18 - 6% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 10 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 92 - 31% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -
Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 4% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 9 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 294 -

```
[[IDc90d5ee43215eb9f82209a490efd81c5]]
= statistics-v1 nfs show-v41
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NFSv41 statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `statistics-v1 nfs show-v41` command displays the following statistics
about the NFSv4.1 operations on each node in the cluster:

* Result of the operations (success or failure)
* Total number of null operations
* Total number of compound operations
* Total number of access operations
* Total number of close operations
* Total number of commit operations
* Total number of create operations
* Total number of delegpurge operations
* Total number of delegret operations
* Total number of getattr operations
* Total number of getfh operations
* Total number of link operations
* Total number of lock operations
* Total number of lockt operations
* Total number of locku operations
* Total number of lookup operations
* Total number of lookupp operations
* Total number of nverify operations
* Total number of open operations
* Total number of openattr operations
* Total number of openconf operations
* Total number of opendowng operations
* Total number of putfh operations
* Total number of putpubfh operations
* Total number of putrootfh operations
* Total number of read operations
* Total number of readdir operations
* Total number of readlink operations
* Total number of remove operations
* Total number of rename operations
* Total number of renew operations
```

```
* Total number of restorefh operations
* Total number of savefh operations
* Total number of secinfo operations
* Total number of setattr operations
* Total number of setcliid operations
* Total number of setcliidconf operations
* Total number of verify operations
* Total number of write operations
* Total number of rellockown operations
* Total number of total operations
* Percent of null operations
* Percent of compound operations
* Percent of access operations
* Percent of close operations
* Percent of commit operations
* Percent of create operations
* Percent of delegpurge operations
* Percent of delegret operations
* Percent of getattr operations
* Percent of getfh operations
* Percent of link operations
* Percent of lock operations
* Percent of lockt operations
* Percent of locku operations
* Percent of lookup operations
* Percent of lookupp operations
* Percent of nverify operations
* Percent of open operations
* Percent of openattr operations
* Percent of openconf operations
* Percent of opendowng operations
* Percent of putfh operations
* Percent of putpubfh operations
* Percent of putrootfh operations
* Percent of read operations
* Percent of readdir operations
* Percent of readlink operations
* Percent of remove operations
* Percent of rename operations
* Percent of renew operations
* Percent of restorefh operations
* Percent of savefh operations
* Percent of secinfo operations
* Percent of setattr operations
* Percent of setcliid operations
* Percent of setCliidconf operations
```

```
* Percent of verify operations
* Percent of write operations
* Percent of rellockown operations

This command is designed to be used to analyze performance characteristics
and to help diagnose issues.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays NFSv4 statistics only
for the specified node.

``[-result {success|failure|all}]`` - Result::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified result (success/failure/all).

``[-null <Counter with Delta>]`` - Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of null operations.

``[-cmpnd <Counter with Delta>]`` - Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of compound operations.

``[-access <Counter with Delta>]`` - Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of access operations.

``[-close <Counter with Delta>]`` - Close Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of close operations.

``[-commit <Counter with Delta>]`` - Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of commit operations.

``[-create <Counter with Delta>]`` - Create Operations::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of create operations.

``[-delpur <Counter with Delta>]`` - Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegpurge operations.

``[-delrtn <Counter with Delta>]`` - Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of delegret operations.

``[-gattr <Counter with Delta>]`` - GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getattr operations.

``[-getfh <Counter with Delta>]`` - GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of getfh operations.

``[-link <Counter with Delta>]`` - Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of link operations.

``[-lock <Counter with Delta>]`` - Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lock operations.

``[-lockt <Counter with Delta>]`` - LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lockt operations.

``[-locku <Counter with Delta>]`` - LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of locku operations.

``[-lookup <Counter with Delta>]`` - Lookup Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookup operations.

``[-lookpp <Counter with Delta>]`` - LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of lookupp operations.

``[-nverfy <Counter with Delta>]`` - Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of nverify operations.
```

```
``[-open <Counter with Delta>]`` - Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of open operations.

``[-opattr <Counter with Delta>]`` - OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openattr operations.

``[-opconf <Counter with Delta>]`` - OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of openconf operations.

``[-opndg <Counter with Delta>]`` - OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of opendowng operations.

``[-putfh <Counter with Delta>]`` - PutFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putfh operations.

``[-putpfh <Counter with Delta>]`` - PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putpubfh operations.

``[-putrfh <Counter with Delta>]`` - PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of putrootfh operations.

``[-read <Counter with Delta>]`` - Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of read operations.

``[-readdr <Counter with Delta>]`` - ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readdir operations.

``[-rlink <Counter with Delta>]`` - ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of readlink operations.

``[-remove <Counter with Delta>]`` - Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of remove operations.

``[-rename <Counter with Delta>]`` - Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rename operations.
```

```
``[-renew <Counter with Delta>]`` - Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of renew operations.

``[-restfh <Counter with Delta>]`` - RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of restorefh operations.

``[-savefh <Counter with Delta>]`` - SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of savefh operations.

``[-secinf <Counter with Delta>]`` - SecInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of secinfo operations.

``[-sattr <Counter with Delta>]`` - SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setattr operations.

``[-sclid <Counter with Delta>]`` - SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliid operations.

``[-scidc <Counter with Delta>]`` - SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of setcliidconf
operations.

``[-verify <Counter with Delta>]`` - Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of verify operations.

``[-write <Counter with Delta>]`` - Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of write operations.

``[-relown <Counter with Delta>]`` - RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of rellockown operations.

``[-total <Counter64 with Delta>]`` - Total Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified number of total nfsv4
operations.
```

```
``[-null-pct <Counter with Delta>]`` - Percent Null Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of null operations.

``[-cmpnd-pct <Counter with Delta>]`` - Percent Compound Procedure::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of compound
operations.

``[-access-pct <Counter with Delta>]`` - Percent Access Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of access operations.

``[-close-pct <Counter with Delta>]`` - Percent Close Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of close operations.

``[-commit-pct <Counter with Delta>]`` - Percent Commit Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of commit operations.

``[-create-pct <Counter with Delta>]`` - Percent Create Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of create operations.

``[-delpur-pct <Counter with Delta>]`` - Percent Delegpurge Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegpurge
operations.

``[-delrtn-pct <Counter with Delta>]`` - Percent Delegret Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of delegret
operations.

``[-gattr-pct <Counter with Delta>]`` - Percent GetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getattr
operations.

``[-getfh-pct <Counter with Delta>]`` - Percent GetFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of getfh operations.

``[-link-pct <Counter with Delta>]`` - Percent Link Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of link operations.
```

```
``[-lock-pct <Counter with Delta>]`` - Percent Lock Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lock operations.

``[-lockt-pct <Counter with Delta>]`` - Percent LockT Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lockt operations.

``[-locku-pct <Counter with Delta>]`` - Percent LockU Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of locku operations.

``[-lookup-pct <Counter with Delta>]`` - Percent Lookup Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookup operations.

``[-lookpp-pct <Counter with Delta>]`` - Percent LookupP Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of lookupp
operations.

``[-nverfy-pct <Counter with Delta>]`` - Percent Nverify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of nverify
operations.

``[-open-pct <Counter with Delta>]`` - Percent Open Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of open operations.

``[-opattr-pct <Counter with Delta>]`` - Percent OpenAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openattr
operations.

``[-opconf-pct <Counter with Delta>]`` - Percent OpenConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of openconf
operations.

``[-opndg-pct <Counter with Delta>]`` - Percent OpenDowng Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of opendowng
operations.

``[-putfh-pct <Counter with Delta>]`` - Percent PutFh Operations::
```

```
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putfh operations.

``[-putpfh-pct <Counter with Delta>]`` - Percent PutPubFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putpubfh
operations.

``[-putrfh-pct <Counter with Delta>]`` - Percent PutRootFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of putrootfh
operations.

``[-read-pct <Counter with Delta>]`` - Percent Read Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of read operations.

``[-readdr-pct <Counter with Delta>]`` - Percent ReadDir Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readdir
operations.

``[-rlink-pct <Counter with Delta>]`` - Percent ReadLink Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of readlink
operations.

``[-remove-pct <Counter with Delta>]`` - Percent Remove Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of remove operations.

``[-rename-pct <Counter with Delta>]`` - Percent Rename Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rename operations.

``[-renew-pct <Counter with Delta>]`` - Percent Renew Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of renew operations.

``[-restfh-pct <Counter with Delta>]`` - Percent RestoreFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of restorefh
operations.

``[-savefh-pct <Counter with Delta>]`` - Percent SaveFh Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of savefh operations.
```

```
``[-secinf-pct <Counter with Delta>]`` - Percent SecInfo Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of secinfo
operations.

``[-sattr-pct <Counter with Delta>]`` - Percent SetAttr Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setattr
operations.

``[-sclid-pct <Counter with Delta>]`` - Percent SetCliId Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliid
operations.

``[-scidc-pct <Counter with Delta>]`` - Percent SetCliIdConf Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of setcliidconf
operations.

``[-verify-pct <Counter with Delta>]`` - Percent Verify Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of verify operations.

``[-write-pct <Counter with Delta>]`` - Percent Write Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of write operations.

``[-relown-pct <Counter with Delta>]`` - Percent RelLockOwn Operations::
If you specify this parameter, the command displays statistics only about
the node or nodes that have the specified percentage of rellockown
operations.

== Examples

The following example displays statistics about the NFSv4.1 operations for
a node named node1:
```

cluster1::> statistics-v1 nfs show-v41 -node node1

Node Value Delta Percent Ops Delta
node1 ----------success----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 6% -
Close Ops: 8 - 3% -

Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 27% -
Getfh Ops: 22 - 8% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 13 - 5% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 8 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 92 - 32% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -
Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 5% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 8 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 286 -
Node Value Delta Percent Ops Delta
node1 ----------failure----------
Null Procs: 0 - 0% -
Cmpnd Procs: 0 - -
Access Ops: 0 - 0% -
Close Ops: 0 - 0% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 0 - 0% -
Getfh Ops: 0 - 0% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 5 - 63% -
Lookupp Ops: 0 - 0% -

Nverify Ops: 0 - 0% -
Open Ops: 2 - 25% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 0 - 0% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 0 - 0% -
Read Ops: 0 - 0% -
Readdir Ops: 0 - 0% -
Readlink Ops: 0 - 0% -
Remove Ops: 0 - 0% -
Rename Ops: 0 - 0% -
Renew Ops: 0 - 0% -
Restorefh Ops: 0 - 0% -
Savefh Ops: 0 - 0% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 1 - 13% -
Setclid Ops: 0 - 0% -
Setclidconf Ops: 0 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 0 - 0% -
Rlockown Ops: 0 - 0% -
Total Ops: 8 -
Node Value Delta Percent Ops Delta
node1 ---------- all ----------
Null Procs: 2 - 1% -
Cmpnd Procs: 92 - -
Access Ops: 16 - 5% -
Close Ops: 8 - 3% -
Commit Ops: 0 - 0% -
Create Ops: 0 - 0% -
Delpur Ops: 0 - 0% -
Delrtn Ops: 0 - 0% -
Getattr Ops: 76 - 26% -
Getfh Ops: 22 - 7% -
Link Ops: 0 - 0% -
Lock Ops: 0 - 0% -
Lockt Ops: 0 - 0% -
Locku Ops: 0 - 0% -
Lookup Ops: 18 - 6% -
Lookupp Ops: 0 - 0% -
Nverify Ops: 0 - 0% -
Open Ops: 10 - 3% -
Openattr Ops: 0 - 0% -
Openconf Ops: 0 - 0% -
Opendowng Ops: 0 - 0% -
Putfh Ops: 92 - 31% -
Putpubfh Ops: 0 - 0% -
Putrootfh Ops: 2 - 1% -
Read Ops: 0 - 0% -
Readdir Ops: 2 - 1% -
Readlink Ops: 0 - 0% -
Remove Ops: 5 - 2% -

Rename Ops: 3 - 1% -
Renew Ops: 0 - 0% -
Restorefh Ops: 11 - 4% -
Savefh Ops: 13 - 4% -
Secinfo Ops: 0 - 0% -
Setattr Ops: 9 - 3% -
Setclid Ops: 1 - 0% -
Setclidconf Ops: 1 - 0% -
Verify Ops: 0 - 0% -
Write Ops: 3 - 1% -
Rlockown Ops: 0 - 0% -
Total Ops: 294 -

```
:leveloffset: -1


= statistics-v1 protocol-request-size commands

:leveloffset: +1


[[IDf97e7bf2176e74eb558e0d0cd619dff6]]
= statistics-v1 protocol-request-size show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display size statistics for CIFS and NFS protocol read and write requests

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays size statistics for CIFS and NFS protocol read and
write requests. The output of the command includes the following
information:

* Node name
* Statistic type
* Average size of request
* Total request count
* Current number of requests in each category of request size
* Number of requests after the command was last executed
```

```
== 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>|local}]`` - Node::
If this parameter is specified, the command displays statistics only for
the specified node.

``[-stat-type <Protocol Type>]`` - RW Request Stat Type::
If this parameter is specified, the command displays only the statistics
of the specified protocol type. Protocol types include the following:
cifs_read, cifs_write, nfs2_read, nfs2_write, nfs3_read, and nfs3_write.

``[-total-req-count <Counter64 with Delta>]`` - Total Request Count::
If this parameter is specified, the command displays only statistics with
the specified total number of requests.

``[-average-size <Counter64 with Delta>]`` - Average Request Size::
If this parameter is specified, the command displays only statistics with
the specified average request size.

``[-histo08 <Counter64 with Delta>]`` - 0 - 511::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo09 <Counter64 with Delta>]`` - 512 - 1023::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo10 <Counter64 with Delta>]`` - 1024 - 2047::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo11 <Counter64 with Delta>]`` - 2048 - 4096::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo12 <Counter64 with Delta>]`` - 4096 - 8191::
If this parameter is specified, the command displays only statistics with
```

```
the specified number of requests in this size range.

``[-histo13 <Counter64 with Delta>]`` - 8192 - 16K::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo14 <Counter64 with Delta>]`` - 16K - 32K::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo15 <Counter64 with Delta>]`` - 32K - 64K::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo16 <Counter64 with Delta>]`` - 64K - 128K::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

``[-histo17 <Counter64 with Delta>]`` - Greater than 128K::
If this parameter is specified, the command displays only statistics with
the specified number of requests in this size range.

== Examples

The following example displays the number of NFS v3 requests in each size
range for only one node in the cluster.
```

cluster1::> statistics protocol-request-size show -stat-type nfs3_* -node node0

Node: node0
Stat Type: nfs3_read
Value Delta
---------- -------- ----------
Average Size: 6 -
Total Request Count:
465947409 -
0-511: 567023 -
512-1023: 4306 -
1K-2047: 175 -
2K-4095: 160404 -
4K-8191: 537576 -
8K-16383: 1742701 -
16K-32767: 1418620 -
32K-65535:
461516604 -
64K-131071: 0 -
128K - : 0 -

Node: node0

Stat Type: nfs3_write
Value Delta
---------- -------- ----------
Average Size: 0 -
Total Request Count:
199294247 -
0-511: 36556 -
512-1023: 3683 -
1K-2047: 745 -
2K-4095: 1413 -
4K-8191: 28643 -
8K-16383:
199223207 -
16K-32767: 0 -
32K-65535: 0 -
64K-131071: 0 -
128K - : 0 -

```
:leveloffset: -1


:leveloffset: -1


= storage-service commands

:leveloffset: +1


[[IDdd0204d7f2ae692840df95bd4ff61aa9]]
= storage-service show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the available storage services

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command displays the available storage services.

NOTE: The available storage services are defined by the type of storage
making up an aggregate.
```

```
== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Selects the available storage services for Vservers that match the
parameter value.

``[-storage-service <text>]`` - Storage Service [.advaced-
priv]#(privilege: advanced)#::
Selects the available storage services whose name matches the parameter
value.

``[-description <text>]`` - Description [.advaced-priv]#(privilege:
advanced)#::
Selects the available storage services whose description matches the
parameter value. This field is a text description of the storage service.

``[-expected-iops-per-tb <integer>]`` - Expected IOPS per TB [.advaced-
priv]#(privilege: advanced)#::
Selects the available storage services whose expected IOPS per TB matches
the parameter value. When multiplied by a number of TB, this field yields
the number of IOPS nominally guaranteed by the storage service. The
multiplier is either the logical used space or the provisioned size of the
storage object, depending on the value of `expected-iops-allocation` .

``[-expected-iops-allocation {used-space|allocated-space}]`` - Expected
IOPS Allocation [.advaced-priv]#(privilege: advanced)#::
Selects the available storage services whose expected IOPS allocation
policy matches the parameter value. The allocation policy is either
``_allocated-space_`` or ``_used-space_`` . When the expected-iops-
allocation policy is ``_allocated-space_`` , the expected IOPS is
calculated based on the size of the storage object. When the expected-
iops-allocation policy is set to ``_used-space_`` , the expected IOPS is
calculated based on the amount of data logically stored in the storage
object.
```

```
``[-peak-iops-per-tb <integer>]`` - Peak IOPS per TB [.advaced-
priv]#(privilege: advanced)#::
Selects the available storage services whose peak IOPS per TB matches the
parameter value. When multiplied by a number of TB, this field yields the
number of IOPS for the maximum Quality of Service (QoS) throttle. The
multiplier is either the logical used space or the provisioned size of the
storage object, depending on the value of `peak-iops-allocation` .

``[-peak-iops-allocation {used-space|allocated-space}]`` - Peak IOPS
Allocation [.advaced-priv]#(privilege: advanced)#::
Selects the available storage services whose peak IOPS allocation policy
matches the parameter value. The allocation policy is either ``_allocated-
space_`` or ``_used-space_`` . When the peak-iops-allocation policy is
``_allocated-space_`` , the peak IOPS is calculated based on the size of
the storage object. When the peak-iops-allocation policy is set to
``_used-space_`` , the peak IOPS is calculated based on the amount of data
logically stored in the storage object.

``[-absolute-min-iops <integer>]`` - Absolute Minimum IOPS [.advaced-
priv]#(privilege: advanced)#::
Selects the available storage services whose absolute minimum IOPS matches
the parameter value. This field is the minimum number of IOPS used as the
Quality of Service (QoS) throttle, if larger than the values calculated
using the IOPS per TB parameters.

``[-target-latency <integer>]`` - Target Latency (ms) [.advaced-
priv]#(privilege: advanced)#::
Selects the available storage service whose target latency matches the
parameter value.

``[-aggr-list <aggregate name>,...]`` - Aggregate List [.advaced-
priv]#(privilege: advanced)#::
Selects the available storage services whose aggregate list matches the
parameter value. The aggregates shown are the only ones used for
provisioning when the corresponding Vserver and storage service are
selected.

== Examples
```

```
cluster1::*> storage-service show
Vserver Storage Service Description
------- -------------- -------------------------------------------
vs1
extreme Extreme Performance
performance Performance
value Value
3 entries were displayed.
```

```
The example above displays all the storage services in the cluster.




:leveloffset: -1


= storage commands

:leveloffset: +1


= storage aggregate commands

:leveloffset: +1


[[ID6d35b5e8975903d53863b00a4464610e]]
= storage aggregate add-disks
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add disks to an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate add-disks` command adds disks to an existing
aggregate. You must specify the number of disks or provide a list of disks
to be added. If you specify the number of disks without providing a list
of disks, the system selects the disks.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the aggregate to which disks are to be added.

``[-diskcount <integer>]`` - Disk Count::
This parameter specifies the number of disks that are to be added to the
aggregate.
```

```
{ ``[-T, -disktype {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD |
VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Disk Type::
This parameter specifies the type of disk that is to be added. It must be
specified with the `-diskcount` parameter when adding disks to a Flash
Pool.
+
Use this parameter when adding spare SSDs to an aggregate to convert it to
a Flash Pool.
+

+
NOTE: Only the aggregates marked as `hybrid-enabled` can be converted to
Flash Pool aggregates. Use xref:{relative_path}storage-aggregate-
modify.html[storage aggregate modify] command to mark the aggregate as
`hybrid-enabled` .
+

+
NOTE: When this parameter is used, disk selection is not influenced by
RAID options `raid.mix.hdd.disktype.capacity` ,
`raid.mix.hdd.disktype.performance` , or `raid.mix.disktype.solid_state` .
Only disks of the specified type are considered eligible for selection.
+


| ``[-C, -diskclass {capacity | performance | archive | solid-state |
array | virtual}]`` - Disk Class::
This parameter specifies the class of disk that is to be added. All disks
that belong to the specified class are considered eligible for selection.
The possible values are:
+
  * capacity = Capacity-oriented, near-line disk types. Includes disk
types FSAS, BSAS and ATA.
  * performance = Performance-oriented, enterprise class disk types.
Includes disk types FCAL and SAS.
  * archive = Archive class SATA disks in multi-disk carrier storage
shelves. Includes disk type MSATA.
  * solid-state = Solid-state drives. Includes disk type SSD, SSD-CAP and
SSD-NVM.
  * array = Logical storage devices backed by storage arrays and used by
ONTAP as disks. Includes disk type LUN.
  * virtual = Virtual disks that are formatted and managed by the
hypervisor. Includes disk type VMDISK.
```

```
+
NOTE: When this parameter is used, disk selection is not influenced by
RAID options `raid.mix.hdd.disktype.capacity` ,
`raid.mix.hdd.disktype.performance` , or `raid.mix.disktype.solid_state` .
+


``[-chksumstyle <aggrChecksumStyle>]`` - Checksum Style::
This parameter specifies the checksum style for the disks to be added to
an aggregate. It is not applicable if `-disklist` or `-mirror-disklist` is
specified. The possible values are ``_block_`` for block checksum and
``_advanced_zoned_`` for advanced zoned checksum (AZCS). By default, disks
with the same checksum style as the aggregate are selected. This behavior
can be overridden by using this parameter to create a mixed checksum
aggregate. A mixed checksum aggregate can support only the ``_block_`` and
``_advanced_zoned_`` checksum styles.

``[-disksize <integer>]`` - Disk Size(GB)::
This parameter specifies the size, in GB, of the disks that are to be
added to the aggregate. Disks with a usable size between 90% and 105% of
the specified size are selected.

| ``[-d, -disklist <disk path name>,...]`` - Disks::
This parameter specifies a list of disks to be added. If you specify the
`-disklist` parameter, you cannot further qualify the list of disks to be
added by count, checksum style, size or type.

``[-mirror-disklist <disk path name>,...]`` - Disks for Mirrored Plex::
This parameter specifies a list of mirror disks to be added. It must
contain the same number of disks specified in `-disklist` parameter. If
you specify the `-mirror-disklist` parameter, you cannot further qualify
the list of disks to be added by count, checksum style or type.

{ ``[-ignore-pool-checks <true>]`` - Don't Enforce Plex Pool Best
Practices::
The disks in a plex are normally required to come from the same SyncMirror
pool. This behavior can be overridden with this parameter when it is set
to ``_true_`` .

``[-f, -allow-mixed-rpm <true>]`` - Allow Disks With Different RPM
Values::
This parameter specifies whether disks that have different RPM values can
be added. For example, SAS disks can rotate at 10,000 or 15,000 RPM. If
this parameter is set to ``_true_`` and a list of disks are provided by
using the `-disklist` parameter, the disks will be added even if the SAS
disks you specify have different RPM values. This parameter works
```

```
similarly for ATA disks, which can rotate at 5,400 or 7,200 RPM.
+
+
NOTE: This parameter is applicable only when the `-disklist` or `-mirror-
disklist` parameter is used.
+


``[-allow-same-carrier <true>]`` - Allow Same RAID Group Within Carrier::
This parameter can be used to allow two disks housed in the same carrier
to be in the same RAID group when you add disks to an aggregate.
+
Having disks in the same carrier in the same RAID group is not desirable
because a carrier failure can cause a simultaneous outage for two disks in
the same RAID group. You can add a disk to an aggregate that causes this
situation, but when an alternate disk becomes available, ONTAP
automatically initiates a series of disk copy operations to put the disks
into different RAID groups. For this reason, you should use this parameter
only when necessary. When possible, allow ONTAP to choose disks that need
to be added to the aggregate.
+
This parameter affects only the add-disks operation. It is not a
persistent attribute of the aggregate.

| ``[-storage-pool <storage pool name>]`` - Storage Pool::
This parameter specifies the name of the SSD storage pool from which
available allocation units are added to a given aggregate. This parameter
cannot be used with the `-disk-list` or `-disk-count` parameters.

``[-allocation-units <integer>]`` - Allocation Units }::
This parameter specifies the number of allocation units to be added to a
given aggregate from an SSD storage pool. Number of allocation units
available and size of each unit can be found using the
xref:{relative_path}storage-pool-show-available-capacity.html[storage pool
show-available-capacity] command. This parameter works only when you also
use the `-storage-pool` parameter.

``[-n, -simulate <true>]`` - Simulate Addition of Disks::
This parameter is used with the disktype and diskcount parameters to
determine which disks would be added without actually performing the
addition of disks operation.

``[-g, -raidgroup {new|all|<raidgroup>}]`` - RAID Group::
This parameter enables the administrator to specify which RAID group will
receive the added disks. If this parameter is not used, the disks are
added to the most recently created RAID group until it is full, then new
```

```
raid groups are created and filled until all the disks are added. If a
RAID group name ``_rgX_`` is specified, the disks are added to that RAID
group. If ``_new_`` is specified, the disks are added to a new RAID group,
even if the disks would fit into an existing RAID group. If ``_all_`` is
specified, the disks are added to existing RAID groups until all existing
RAID groups are full. Then ONTAP creates one or more new RAID groups and
adds the remaining disks to the new groups. If the disk type or checksum
style parameters are specified with this parameter, the command operates
only on the RAID groups with the matching disk type or checksum style,
even if ``_all_`` is specified.

``[-cache-raid-group-size <integer>]`` - RAID Group Size for Cache Tier::
This parameter specifies the maximum number of disks that can be included
in an SSD RAID group for this aggregate.
+
NOTE: This parameter is applicable only when adding SSDs for the first
time to a hybrid-enabled aggregate. If this parameter is not used when the
first SSDs are added to the aggregate, the maximum RAID group size for the
SSD cache is set to the default SSD RAID group size for the RAID type of
the SSD cache.
+


``[-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}]`` - RAID Type::
This parameter specifies the type for the new RAID groups that would be
created while adding disks to the aggregate. Use this parameter when you
add the first RAID group comprised of SSDs to a hybrid-enabled aggregate.
The values are ``_raid4_`` for RAID4, ``_raid_dp_`` for RAID Double
Parity, and ``_raid_tec_`` for RAID-TEC. The default value is the type of
RAID groups of the aggregate, except for RAID-TEC hybrid-enabled
aggregates where the SSD tier will default to ``_raid_dp_`` . An aggregate
might include a mix of different RAID types.

== Examples

The following example adds 10 disks to an aggregate named aggr0. The disks
are added to a RAID group named rg1:
```

cluster1::> storage aggregate add-disks -aggregate aggr0 -diskcount 10 -raidgroup rg1

```
In this example, an aggregate is converted to a Flash Pool aggregate using
SSD capacity from a storage pool. The aggregate was created using RAID-DP
for the hard disks and the SSDs are added using RAID4.
```

cluster1::> storage aggregate add-disks -aggregate FlashPool -storage-pool SP1 -allocation-units 1 -raidtype

```
== Related Links

* xref:{relative_path}storage-aggregate-modify.html[storage aggregate
modify]
* xref:{relative_path}storage-pool-show-available-capacity.html[storage
pool show-available-capacity]




[[ID9eb801c2767c15db22e638c385151072]]
= storage aggregate auto-provision
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Recommend and create new aggregates in the cluster

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command analyzes available spare disks in the cluster, and it
provides a recommendation how spare disks should be used to create
aggregates according to best practices. The command prints the summary of
recommended aggregates including their names and usable size. It then
prompts the user whether or not the aggregates should be created as
recommended. On positive response, ONTAP creates the aggregates as
described in the recommendation.



The command parameters allow to restrict the command to some nodes in the
cluster, print more details about recommended aggregates, and to skip the
prompt.

== Parameters

``[-nodes {<nodename>|local}]`` - List of Nodes::
Comma separated list of node names to which the command applies. If this
parameter is not used, the command applies to all nodes in the cluster.
```

```
``[-verbose <true>]`` - Report More Details::
Report additional details about recommended aggregates and spare disks.
Per node summary shows number and total size of aggregates to create,
discovered spares, and also remaining spare disks and partitions after
aggregate creation. RAID group layout shows how spare disks and partitions
will be used in new data aggregates to be created. The last table shows
spare disks and partitions remaining unused after aggregate creation.

``[-skip-confirmation <true>]`` - Skip the Confirmation and Create
Recommended Aggregates::
When this parameter is used, the command automatically creates the
recommended aggregates. When this parameter is not used, the command
checks to proceed with aggregate creation or not.
+
NOTE: The command is not affected by the CLI session setting:
xref:{relative_path}set.html[set]`-confirmations```_on_`` /``_off_`` .
+


== Examples
```

cluster1::storage aggregate> auto-provision
Node New Data Aggregate Usable Size
--- --- ---
node1 node1_SSD_1 3.66TB
node2 node2_SSD_1 3.66TB
--- --- ---
Total: 2 new data aggregates 7.32TB

Do you want to create recommended aggregates? {y|n}: n

cluster1::storage aggregate> auto-provision -verbose
Per node summary of new aggregates to create, discovered spares, and also
remaining spare disks and partitions after aggregate creation:
New Total New -Discovered Spare- -Remaining Spare-
Node Aggrs Usable Size Disks Partitions Disks Partitions
--- --- --- --- --- --- ---
node1 1 3.66TB 6 48 1 0
node2 1 3.66TB 6 48 1 0
--- --- --- --- --- --- ---
Total: 2 7.32TB 12 96 2 0

New data aggregates to create with counts of
disks and partitions to be used:
Is Usable -Devices To Use-
Node New Data Aggregate Mirrored Size Disks Partitions
--- --- --- --- --- ---
node1 node1_SSD_1 false 3.66TB 5 48
node2 node2_SSD_1 false 3.66TB 5 48

RAID group layout showing how spare disks and partitions will be used
in new data aggregates to be created:

RAID Group In New Disk Usable Disk Or ---Count---
Data Aggregate To Be Created Type Size Partition Data Parity
--- --- --- --- --- ---
/node1_SSD_1/plex0/rg0 SSD 81.97GB partition 22 2
/node1_SSD_1/plex0/rg1 SSD 81.97GB partition 22 2
/node1_SSD_1/plex0/rg2 SSD 185.5GB disk 3 2
/node2_SSD_1/plex0/rg0 SSD 81.97GB partition 22 2
/node2_SSD_1/plex0/rg1 SSD 81.97GB partition 22 2
/node2_SSD_1/plex0/rg2 SSD 185.5GB disk 3 2

Details about spare disks and partitions remaining after aggregate creation:
Disk Device Disk Or Pool Remaining
Node Type Usable Size Partition Number Spares
--- --- --- --- --- ---
node1 SSD 185.5GB disk Pool0 1
node2 SSD 185.5GB disk Pool0 1

Do you want to create recommended aggregates? {y|n}: y

Info: Aggregate auto provision has started. Use the "storage aggregate
show-auto-provision-progress" command to track the progress.

```
== Related Links

* xref:{relative_path}set.html[set]




[[ID8ee30549053ead1f6a8340d95c3a617e]]
= storage aggregate create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate create` command creates an aggregate. An aggregate
consists of disks. You must specify the number of disks or provide a list
```

```
of disks to be added to the new aggregate. If you specify the number of
disks without providing a list of disks, the system selects the disks.

When creating an aggregate, you can optionally specify the aggregate's
home node, the RAID type for RAID groups on the aggregate, and the maximum
number of disks that can be included in a RAID group.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the name of the aggregate that is to be created.

``[-chksumstyle <aggrChecksumStyle>]`` - Checksum Style::
This parameter specifies the checksum style for the aggregate. The values
are ``_block_`` for Block Checksum and ``_advanced_zoned_`` for Advanced
Zoned Checksum (AZCS).

``-diskcount <integer>`` - Number Of Disks::
This parameter specifies the number of disks that are to be included in
the aggregate, including the parity disks. The disks in this newly created
aggregate come from the pool of spare disks. The smallest disks in this
pool are added to the aggregate first, unless you specify the `-disksize`
parameter.

``[-R, -diskrpm <integer>]`` - Disk RPM::
This parameter specifies the RPM of the disks on which the aggregate is to
be created. The possible values include 5400, 7200, 10000, and 15000.

``[-disksize <integer>]`` - Disk Size(GB)::
This parameter specifies the size, in GB, of the disks on which the
aggregate is to be created. Disks with a usable size between 90% and 105%
of the specified size are selected.

{ ``[-T, -disktype {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD |
VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Disk Type::
This parameter specifies the type of disk on which the aggregate is to be
created.
+
NOTE: When this parameter is used, disk selection is not influenced by
RAID options `raid.mix.hdd.disktype.capacity` ,
`raid.mix.hdd.disktype.performance` , or `raid.mix.disktype.solid_state` .
Only disks of the specified type are considered eligible for selection.
+



| ``[-C, -diskclass {capacity | performance | archive | solid-state |
array | virtual}]`` - Disk Class::
```

```
This parameter specifies the class of disks on which the aggregate is to
be created. All disks that belong to the specified class are considered
eligible for selection. The possible values are:
+
  * capacity = Capacity-oriented, near-line disk types. Includes disk
types FSAS, BSAS and ATA.
  * performance = Performance-oriented, enterprise class disk types.
Includes disk types FCAL and SAS.
  * archive = Archive class SATA disks in multi-disk carrier storage
shelves. Includes disk type MSATA.
  * solid-state = Solid-state drives. Includes disk type SSD, SSD-CAP and
SSD-NVM.
  * array = Logical storage devices backed by storage arrays and used by
ONTAP as disks. Includes disk type LUN.
  * virtual = Virtual disks that are formatted and managed by the
hypervisor. Includes disk type VMDISK.



+
NOTE: When this parameter is used, disk selection is not influenced by
RAID options `raid.mix.hdd.disktype.capacity` ,
`raid.mix.hdd.disktype.performance` , or `raid.mix.disktype.solid_state` .
+



``[-m, -mirror <true>]`` - Is Mirrored::
This parameter specifies that the new aggregate be mirrored (have two
plexes). If this parameter is set to ``_true_`` , the specified disks are
split between the two plexes. By default, the new aggregate will not be
mirrored. You cannot use the `-mirror` parameter when supplying a specific
list of disks with either the `-disklist` or `-mirror-disklist`
parameters.

``[-pool <aggrSparePool>]`` - Spare Pool::
This parameter specifies the SyncMirror pool to be used to supply the
disks for the aggregate. Valid values are Pool0 or Pool1.

| ``-d, -disklist <disk path name>,...`` - Disks for First Plex::
This parameter specifies a list of disks to be added to the new aggregate.
If you specify the `-disklist` parameter, you cannot further qualify the
list of disks to be added by count, checksum style, type, size, or RPM.
You cannot use the `-disklist` parameter when the `-mirror` parameter is
set to ``_true_`` .

``[-mirror-disklist <disk path name>,...]`` - Disks for Mirrored Plex::
This parameter specifies a list of mirror disks to be added to the new
```

```
mirrored aggregate. It must contain the same number of disks specified in
`-disklist` parameter. If you specify the `-mirror-disklist` parameter,
you cannot further qualify the list of disks to be added by count,
checksum style, type, size, or RPM. You cannot use the `-mirror-disklist`
parameter when the `-mirror` parameter is set to ``_true_`` .

``[-ignore-pool-checks <true>]`` - Don't Enforce Plex Pool Best
Practices::
The disks in a plex are normally required to come from the same SyncMirror
pool. This behavior can be overridden with this parameter when it is set
to ``_true_`` . This option cannot be used when the `-mirror` option is
set to ``_true_``

``[-f, -allow-mixed-rpm <true>]`` - Allow Disks With Different RPM
Values::
This parameter specifies whether the aggregate can contain disks that have
different RPM values. For example, SAS disks can rotate at 10,000 or
15,000 RPM. If this parameter is set to ``_true_`` and a list of disks are
provided by using the `-disklist` parameter, the aggregate will be created
even if the SAS disks you specify have different RPM values. This
parameter works similarly for ATA disks, which can rotate at 5,400 or
7,200 RPM.

``[-allow-same-carrier <true>]`` - Allow Same RAID Group Within Carrier::
This parameter can be used to allow two disks housed in the same carrier
to be in the same RAID group when you add disks to an aggregate.
+
Having disks in the same carrier in the same RAID group is not desirable
because a carrier failure can cause a simultaneous outage for two disks in
the same RAID group. You create an aggregate with this characteristic, but
when an alternate disk becomes available, ONTAP automatically initiates a
series of disk copy operations to put the disks into different RAID
groups. For this reason, you should use this parameter only when
necessary. When possible, allow ONTAP to choose the disks from which to
create the aggregate.
+
This parameter affects only the aggregate creation operation. It is not a
persistent attribute of the aggregate.

``[-node {<nodename>|local}]`` - Node::
This parameter specifies the home node for the aggregate. If this
parameter is not specified, ONTAP selects the node where the aggregate is
created.

``-ha-policy {sfo|cfo}`` - HA Policy::
This parameter specifies the high-availability policy of the aggregate.
```

```
``[-s, -maxraidsize <integer>]`` - Max RAID Size::
This parameter specifies the maximum number of disks that can be included
in a RAID group.

``[-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}]`` - RAID Type::
This parameter specifies the type for RAID groups on the aggregate. The
values are ``_raid4_`` for RAID4, ``_raid_dp_`` for RAID Double Parity,
and ``_raid_tec_`` for RAID Triple-Erasure-Code. The default setting is
``_raid_dp_`` unless the disks are HDDs with a capacity larger than 4 TB,
in which case the default will be ``_raid_tec_`` . This parameter is not
needed for array LUNs because they are always created with the ``_raid0_``
raidtype. ``_raid4_`` is not compatible with shared disks unless the
shared disks belong to a storage pool.

``[-simulate <true>]`` - Simulate Aggregate Provisioning Operation::
This option simulates the aggregate creation and prints the layout of the
new aggregate.

``[-force-small-aggregate <true>]`` - Force the Creation of a Small
Aggregate [.advaced-priv]#(privilege: advanced)#::
This parameter can be used to force the creation of a 2-disk RAID4
aggregates, or a 3-disk or 4-disk RAID-DP aggregate.

``[-is-autobalance-eligible {true|false}]`` - Is Eligible for Auto Balance
Aggregate [.advaced-priv]#(privilege: advanced)#::
This specifies whether the aggregate will be considered by the Auto
Balance Aggregate feature. If the Auto Balance Aggregate feature is not
used, this field is not used. When this parameter is set to ``_true_`` the
Auto Balance Aggregate feature might recommend moving volumes to or from
this aggregate in order to balance system workload. When this parameter is
set to ``_false_`` the aggregate will not be considered as a destination
for the Auto Balance Aggregate feature allowing for predictability in data
placement. The default value is ``_false_`` .

``[-L, -snaplock-type {non-snaplock|compliance|enterprise}]`` - SnapLock
Type::
This parameter specifies the type of SnapLock aggregate to be created. In
order to create a SnapLock Compliance aggregate, specify ``_compliance_``
. To create a SnapLock Enterprise aggregate, specify ``_enterprise_`` .

``[-autobalance-unbalanced-threshold-percent <integer>]`` - Threshold When
Aggregate Is Considered Unbalanced (%) [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the space used threshold percentage that will
cause the Auto Balance Aggregate feature to consider an aggregate as
unbalanced.
```

```
``[-autobalance-available-threshold-percent <integer>]`` - Threshold When
Aggregate Is Considered Balanced (%) [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the threshold percentage which will determine if
an aggregate is a target destination for a move. The Auto Balance
Aggregate feature will attempt to move volumes from an unbalanced
aggregate until it is under this percentage.

``[-encrypt-with-aggr-key {true|false}]`` - Enable Aggregate level
Encryption::
This parameter specifies the data encryption policy for the contained
volumes. If this parameter is set to ``_true_`` , then by default, the
volumes created in this aggregate will be encrypted, using the aggregate
level encryption keys.

== Examples

The following example creates an aggregate named aggr0 on a home node
named node0. The aggregate contains 20 disks and uses RAID-DP. The
aggregate contains regular FlexVol volumes:
```

cluster1::> storage aggregate create -aggregate aggr0 -node node0
-diskcount 20 -raidtype raid_dp -volume-style flex

```
The following example creates an aggregate named aggr0 on a home node
named node0. The aggregate contains the disks specified and uses RAID-DP
```

cluster1::> storage aggregate create -aggregate aggr0 -node node0
-disklist 1.0.15,1.0.16,1.0.17,1.0.18,1.0.19 -raidtype raid_dp

```
The following example creates an aggregate named aggr0 on a home node
named node0. The aggregate contains 20 disks of size 6 TB and of type FSAS
and uses RAID-TEC:
```

cluster1::> storage aggregate create -aggregate aggr0 -node node0
-diskcount 20 -raidtype raid_tec -disksize 6000 -disktype FSAS

```
The following example creates a mirrored aggregate named aggr0 on the
local node. The aggregate contains 10 disks in each plex:
```

cluster1::> storage aggregate create -aggregate aggr0 -mirror
-diskcount 20

```
[[ID4a77573d8b7d549eed85c8a9db8c6653]]
= storage aggregate delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate delete` command deletes a storage aggregate. The
command fails if there are volumes present on the aggregate. If the
aggregate has an object store attached to it, then in addition to deleting
the aggregate the command deletes the objects in the object store as well.
No changes are made to the object store configuration as part of this
command.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the aggregate that is to be deleted.

``[-preserve-config-data <true>]`` - Delete Physical Aggregate but
Preserve Configuration Data [.advaced-priv]#(privilege: advanced)#::
Deletes the physical aggregate, but preserves the aggregate configuration
data. The aggregate must not have any disks associated with it. If the
parameter `-preserve-config-data` is specified without a value, the
default value is ``_true_`` ; if this parameter is not specified, the
default value is ``_false_`` .

== Examples

The following example deletes an aggregate named aggr1:
```

cluster1::> storage aggregate delete -aggregate aggr1

```
[[ID59fb7ad9fc0b1f2b4ea7bb06764e9559]]
= storage aggregate mirror
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Mirror an existing aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate mirror` command adds a plex to an existing
unmirrored aggregate. You can specify a list of disks to be used for the
mirrored plex. If you do not specify the disks, the system automatically
selects the disks based on the aggregate's existing plex.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the aggregate to mirror.

``[-f, -allow-mixed-rpm <true>]`` - Allow Disks With Different RPM
Values::
This parameter specifies whether disks that have different RPM values can
be used. For example, SAS disks can rotate at 10,000 or 15,000 RPM. If
this parameter is set to ``_true_`` and a list of disks are provided by
using the `-mirror-disklist` parameter , the disks will be added even if
the SAS disks you specify have different RPM values. This parameter works
similarly for ATA disks, which can rotate at 5,400 or 7,200 RPM.
+
+
NOTE: This parameter is only applicable when the `-mirror-disklist`
parameter is used.
+



``[-d, -mirror-disklist <disk path name>,...]`` - Disks for Mirrored
Plex::
This parameter specifies a list of disks to be used for the plex to be
added. It must contain the same number of disks as the existing plex of
the unmirrored aggregate specified using the `-aggregate` parameter.

``[-ignore-pool-checks <true>]`` - Don't Enforce Plex Pool Best
Practices::
For maximum reliability, all disks from a plex should come from the same
SyncMirror pool, and the disks for the second plex should all come from
the other pool. If needed, this behavior can be overridden by setting this
```

```
parameter to ``_true_`` . This parameter can be used only with the `-
mirror-disklist` parameter .

``[-f, -allow-same-carrier <true>]`` - Allow Same RAID Group Within
Carrier::
This parameter can be used to allow two disks housed in the same carrier
to be in the same RAID group for a mirrored aggregate. Having disks in the
same carrier in the same RAID group is not desirable, because a carrier
failure can cause a simultaneous outage for two disks in the same RAID
group. For this reason, this configuration is not allowed by default. This
restriction can be overridden by setting this parameter to ``_true_`` .
+
+
NOTE: This parameter is accepted only when the `-mirror-disklist`
parameter is used.
+


``[-n, -simulate <true>]`` - Simulate Mirroring of an Existing Aggregate::
This option simulates the mirroring of an existing aggregate and prints
the layout of the new plex.

== Examples

The following example mirrors an unmirrored aggregate aggr1:
```

cluster1::> storage aggregate mirror -aggregate aggr1

```
The following example mirrors an unmirrored aggregate aggr1. The specified
disks are used for the new plex.
```

cluster1::> storage aggregate mirror -aggregate aggr1 -mirror-disklist 1.2.12, 1.2.14, 1.2.16

```
[[ID2b8c4011806b0d15de978bdb467ad535]]
= storage aggregate modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify aggregate attributes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `storage aggregate modify` command can be used to modify attributes of
an aggregate such as RAID type and maximum RAID group size.

Changing the RAID type immediately changes the RAID group type for all
RAID groups in the aggregate.

Changing the maximum RAID size does not cause existing RAID groups to grow
or to shrink; rather, it affects the size of RAID groups created in the
future, and determines whether more disks can be added to the RAID group
that was most recently created.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the storage aggregate that is to be modified.

``[-T, -disktype {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD |
VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Disk Type::
This parameter specifies the disk type of the RAID groups to be modified.
In case of Flash Pool, it specifies either the HDD tier or the SSD tier.
If the HDD tier is composed of more than one type of disk, specifying any
of the disk types in use causes that tier to be modified. If the current
aggregate RAID type is mixed_raid_type, this parameter is mandatory.

``[-free-space-realloc {on|off}]`` - Free Space Reallocation::
This parameter specifies whether free space reallocation is enabled on the
aggregate.
+
Free space reallocation optimizes the free space in an aggregate
immediately before ONTAP writes data to the blocks in that aggregate.
+
The default setting is `off.`

``[-ha-policy {sfo|cfo}]`` - HA Policy::
This parameter specifies the high-availability policy to be used in the
context of a root recovery procedure. Do not modify this setting unless
directed to do so by a customer support representative.

``[-percent-snapshot-space <percent>]`` - Space Reserved for Snapshot
Copies::
This parameter is used to set the space reserved for snapshots to the
specified value. For example, to set the snapshot reserve to 5%, you
should enter `-percent-snapshot-space```_5_`` .
```

```
``[-space-nearly-full-threshold-percent <percent>]`` - Aggregate Nearly
Full Threshold Percent::
This optionally specifies the percentage at which the aggregate is
considered nearly full, and above which an EMS warning will be generated.
The default value is 95%. The maximum value for this option is 99%.
Setting this threshold to 0 disables the aggregate nearly full space
alerts.

``[-space-full-threshold-percent <percent>]`` - Aggregate Full Threshold
Percent::
This optionally specifies the percentage at which the aggregate is
considered full, and above which a critical EMS error will be generated.
The default value is 98%. The maximum value for this option is 100%.
Setting this threshold to 0 disables the aggregate full space alerts.

``[-hybrid-enabled {true|false}]`` - Hybrid Enabled::
If the hybrid-enabled option is set to "true", the aggregate is marked as
hybrid_enabled, that is, the aggregate can contain a mix of SSDs and HDDs
(Hard Disk Drives, e.g., SAS, SATA, and/or FC). By default, aggregates
cannot be marked "hybrid_enabled" if the aggregate contains FlexVol
volumes that cannot be write cached. A FlexVol cannot be write-cached if
it is part of an aggregate created in ONTAP 7. Use `-force-hybrid-enabled`
to over-ride this behavior.

``[-f, -force-hybrid-enabled <true>]`` - Force Marking of Aggregate as
Hybrid Enabled::
By default, aggregates cannot be marked "hybrid_enabled" if the aggregate
contains FlexVol volumes that cannot be write cached. A FlexVol cannot be
write-cached if it is part of an aggregate created in ONTAP 7. Use `-
force-hybrid-enabled` to over-ride this behavior. Note that read caching
will be enabled on these FlexVol volumes, but write caching will be
disabled.

``[-s, -maxraidsize <integer>]`` - Max RAID Size::
This parameter specifies the maximum number of disks that can be included
in a RAID group for this aggregate.
+
NOTE: For Flash Pool aggregates, this option controls the maximum size of
the HDD RAID groups.
+


``[-cache-raid-group-size <integer>]`` - Flash Pool SSD Tier Maximum RAID
Group Size::
This parameter specifies the maximum number of disks that can be included
in a SSD RAID group for this Flash Pool.
```

```
+
NOTE: This parameter is applicable only for Flash Pool aggregates.
+


``[-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}]`` - RAID Type::
This parameter specifies the RAID type for RAID groups on the aggregate.
The possible values are ``_raid4_`` for RAID4, ``_raid_dp_`` for RAID-DP,
and ``_raid_tec_`` for RAID-TEC. If you change the RAID type from RAID4 to
RAID-DP, each RAID group allocates a spare disk for the group's second
parity disk and begins a reconstruction process. If you change the RAID
type from RAID-DP to RAID-TEC, each RAID group allocates a spare disk for
the group's third parity disk and begins a reconstruction process.
Changing the RAID type from RAID4 to RAID-TEC or vice-versa is not
supported. To change the RAID type from RAID4 to RAID-TEC, first change
from RAID4 to RAID-DP and then to RAID-TEC.

``[-resyncsnaptime <integer>]`` - SyncMirror Resync Snapshot Frequency in
Minutes::
This parameter sets the mirror resynchronization snapshot frequency to be
the given number of minutes. The default value is 5 (minutes).

``[-state <aggregate state>]`` - State::
This deprecated parameter specifies the state of the aggregate. The
possible values are as follows:
+
  * online - Immediately sets the aggregate online. All volumes on the
aggregate are set to the state they were in when the aggregate was taken
offline or restricted. The preferred command to bring an aggregate online
is `storage aggregate online.`
  * offline - Takes an aggregate offline. You cannot take an aggregate
offline if any of its volumes are online. The preferred command to take an
aggregate offline is `storage aggregate offline.`
  * restricted - Restricts the aggregate. You cannot restrict an aggregate
if any of its volumes are online. The preferred command to restrict an
aggregate is `storage aggregate restrict.`



``[-is-autobalance-eligible {true|false}]`` - Is Eligible for Auto Balance
Aggregate [.advaced-priv]#(privilege: advanced)#::
This specifies whether the aggregate is considered by the Auto Balance
Aggregate feature. If the Auto Balance Aggregate feature is not used, this
field is not used. When this parameter is set to ``_true_`` the Auto
Balance Aggregate feature might recommend moving volumes to or from this
aggregate in order to balance system workload. When this parameter is set
```

```
to ``_false_`` the aggregate will not be considered as a destination for
the Auto Balance Aggregate feature allowing for predictability in data
placement. The default value is ``_false_`` .

``[-autobalance-unbalanced-threshold-percent <integer>]`` - Threshold When
Aggregate Is Considered Unbalanced (%) [.advaced-priv]#(privilege:
advanced)#::
This parameter sets the space used threshold percentage that will cause
the Auto Balance Aggregate feature to consider an aggregate as unbalanced.

``[-autobalance-available-threshold-percent <integer>]`` - Threshold When
Aggregate Is Considered Balanced (%) [.advaced-priv]#(privilege:
advanced)#::
This parameter sets the threshold percentage which will determine if an
aggregate is a target destination for a move. The Auto Balance Aggregate
feature will attempt to move volumes from an unbalanced aggregate until it
is under this percentage.

``[-resync-priority {high(fixed)|high|medium|low}]`` - Resynchronization
Priority::
This parameter specifies the new resynchronization priority value for the
specified aggregate. This field cannot be modified for unmirrored or ONTAP
system aggregates.
+
Possible values for this parameter are:
+

+
  * high: Mirrored data aggregates with this priority value start
resynchronization first.
  * medium: Mirrored data aggregates with this priority value start
resynchronization after all the system aggregates and data aggregates with
'high' priority value have started resynchronization.
  * low: Mirrored data aggregates with this priority value start
resynchronization only after all the other aggregates have started
resynchronization.



``[-single-instance-data-logging {off|on}]`` - Enable SIDL::
This parameter specifies whether Single Instance Data Logging feature is
enabled on the aggregate and the constituent volumes on the aggregate.
This feature improves user write performance by optimizing the amount of
data nvlogged by user writes on platforms where NVRAM and secondary
storage are of same media type.
```

```
``[-is-inactive-data-reporting-enabled {true|false}]`` - Inactive Data
Reporting Enabled::
This parameter specified whether the reporting of how much user data is
inactive should be enabled on the aggregate and volumes on the aggregate.
This parameter is not allowed on FabricPools.

``[-encrypt-with-aggr-key {true|false}]`` - Enable Aggregate level
Encryption::
This parameter specifies that the volumes within the new aggregate can be
encrypted with aggregate keys. If this parameter is set to ``_true_`` ,
the aggregate will support encryption with aggregate keys.

``[-force-disable-encrypt-with-aggr-key <true>]`` - Force disable NAE.
Skip aggregate snapshot check.::
This parameter allows disabling NetApp Aggregate Encryption (NAE) on an
aggregate if the user is certain there is no aggregate snapshot for that
aggregate containing NAE volumes. If the parameter is set to ``_true_`` ,
aggregate snapshot check is skipped and NAE is disabled.

``[-azcs-read-optimization {on|off}]`` - azcs read optimization::
This parameter specifies whether azcs read optimization is enabled on the
aggregate. This feature improves read performance on cloud platforms.

== Examples

The following example changes all RAID groups on an aggregate named aggr0
to use RAID-DP:
```

cluster1::> storage aggregate modify -aggregate aggr0 -raidtype raid_dp

```
The following example changes all RAID groups with FSAS disks in an
aggregate named aggr0 to use RAID-TEC:
```

cluster1::> storage aggregate modify -aggregate aggr0 -disktype FSAS -raidtype raid_tec

```
[[IDbd6b1dda1654673977ad82d796dc8da8]]
= storage aggregate offline
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Offline an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate offline` command takes an aggregate offline.

If you are taking a root aggregate offline, the node owning the aggregate
must be in maintenance mode.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
The name of the aggregate to be taken offline.

== Examples

The following example takes an aggregate named aggr1 offline:
```

cluster1::> storage aggregate offline -aggregate aggr1

```
[[IDfe6018c29a8d28fb2fb5bd437ed09430]]
= storage aggregate online
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Online an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate online` command brings an aggregate online if the
aggregate is in offline or restricted state. If an aggregate is in an
inconsistent state, it must be brought to a consistent state before it can
be brought online. If you have an aggregate that is in an inconsistent
state, contact technical support.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
The name of the aggregate to be brought online.

== Examples

The following example brings an aggregate named aggr1 online:
```

cluster1::> storage aggregate online -aggregate aggr1

```
[[ID464f83061f28c08cf46b5b52bdefe608]]
= storage aggregate remove-stale-record
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a stale aggregate record

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate remove-stale-record` command removes a stale
storage aggregate record on disk. A stale aggregate record refers to an
aggregate that has been removed from the storage system, but whose
information remains recorded on disk. Stale aggregate records are
displayed in the nodeshell aggr status -r command, but the storage
aggregate show command does not show the aggregate as hosted on that node.

== Parameters

``-aggregate <aggregate name>`` - Aggregate [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the aggregate that corresponds to the stale
aggregate record that is to be deleted.

``-nodename {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node that contains the aggregate.

== Examples

The following example removes a stale aggregate record that refers to
aggregate "aggr1":
```

cluster1::> storage aggregate remove-stale-record -aggregate aggr1 -nodename node1

```
[[IDb05c94610659dd9bd02f8a7bb04fdc6b]]
= storage aggregate rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate rename` command renames an aggregate.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the aggregate to be renamed.

``-newname <aggregate name>`` - New Name::
This parameter specifies the new name for the aggregate.

== Examples

The following example renames an aggregate named aggr5 as sales-aggr:
```

cluster1::> storage aggregate rename -aggregate aggr5 -newname sales-aggr

```
[[IDa2daf8a6d5acee4f8f70c0ae893cd7a9]]
= storage aggregate restrict
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restrict an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate restrict` command puts an aggregate in restricted
state to make data in the aggregate's volumes unavailable to clients. When
an aggregate is in restricted state data access is not allowed. However,
few operations such as aggregate copy, parity recomputation, scrub and
RAID reconstruction are allowed. You can also use this command if you want
the aggregate to be the target of an aggregate copy or SnapMirror
replication operation.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
The name of the aggregate to be restricted.

== Examples

The following example restricts an aggregate named aggr1:
```

cluster1::> storage aggregate restrict -aggregate aggr1

```
[[ID40e200d8157b931a32ae4ffc44cdf87f]]
= storage aggregate scrub
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Aggregate parity scrubbing

*Availability:* This command is available to _cluster_ administrators at
```

```
the _admin_ privilege level.

== Description

The `storage aggregate scrub` command scrubs an aggregate for media and
parity errors. Parity scrubbing compares the data disks to the parity
disks in their RAID group and corrects the parity disks contents, as
required. If no name is given, parity scrubbing is started on all online
aggregates.

NOTE: By default, scrubs are scheduled to run for a specified time on a
weekly basis. However, you can use this command to run scrubs manually to
check for errors and data inconsistencies.



== Parameters

{ ``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the aggregate to be scrubbed for errors.

``[-plex <text>]`` - Plex::
This parameter specifies the name of the plex to scrub. If this parameter
is not specified, the command scrubs the entire aggregate.

``[-raidgroup <text>]`` - RAID Group::
This parameter specifies the RAID group to be scrubbed. If this parameter
is not specified, the command scrubs the entire aggregate.
+
+
NOTE: This parameter is only applicable when the `-plex` parameter is
used.
+



| ``-node {<nodename>|local}`` - Node }::
This parameter specifies the name of the node associated with the
aggregate to be scrubbed. The value ``_local_`` specifies the current
node.

``-action {start|stop|resume|suspend|status}`` - Action::
This parameter specifies the action to be taken. The possible actions are:
+
  * start - Starts a scrub.
  * stop - Permanently stops a scrub. A stopped scrub cannot be resumed.
  * resume - Resumes a suspended parity scrub.
  * suspend - Suspends a parity scrub.
```

```
  * status - Displays the current status of a scrub.



== Examples

The following example starts a scrub on a RAID group named rg0 of plex
named plex0 on an aggregate named aggr0:
```

cluster1::> storage aggregate scrub -aggregate aggr0 -raidgroup rg0 -plex plex0 -action start

```
The following example queries the status of a scrub:
```

cluster1::> storage aggregate scrub -aggregate aggr0 -raidgroup rg0 -plex plex0 -action status

Raid Group:/aggr0/plex0/rg0, Is Suspended:false, Last Scrub:Sun Nov 13
01:30:55 2011
, Percentage Completed:7%

```
The following example starts a scrub on plex1 of an aggregate named aggr1:
```

cluster1::> storage aggregate scrub -aggregate aggr1 -plex plex1 -action start

```
The following example queries the status of plex1 of an aggregate named
aggr1:
```

cluster1::> storage aggregate scrub -aggregate aggr1 -plex plex1 -action status

Raid Group:/aggr1/plex1/rg0, Is Suspended:false, Last Scrub:Sun Nov 13 02:07:29
2011
, Percentage Completed:1%

```
The following example queries the status of all the plexes for an
aggregate named aggr1:
```

cluster1::> storage aggregate scrub -aggregate aggr1 -action status

Raid Group:/aggr1/plex0/rg0, Is Suspended:false, Last Scrub:Sun Nov 13 01:58:06
2011

Raid Group:/aggr1/plex1/rg0, Is Suspended:false, Last Scrub:Sun Nov 13 02:07:29
2011
, Percentage Completed:4%

```
[[ID1999c661c60cc0ef402a39e2f903976b]]
= storage aggregate show-auto-provision-progress
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate auto provision status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show-auto-provision-progress` command displays the
status of the most recent auto provision operation. The command output
displays the progress for all the aggregates included in the provisioning
operation. The command displays the following information about each
aggregate:

* Aggregate
* Current Usable Size
* Target Usable Size
* Provisioning Progress



== Examples

The following example displays the information about all aggregates that
are provisioned during the aggregate auto provision operation:
```

cluster1::> aggr auto-provision
Node New Data Aggregate Usable Size
--- --- ---
node1 node1_SSD_1 23.65GB
node2 node2_SSD_1 23.65GB
--- --- ---
Total: 2 new data aggregates 47.30GB

Do you want to create recommended aggregates? {y|n}: y

Info: Aggregate auto provision has started. Use the "storage aggregate
show-auto-provision-progress" command to track the progress.

cluster1::> storage aggregate show-auto-provision-progress

```
Current Target
Aggregate Usable Size Usable Size Provisioning Progress
--------------------------- ------------ ------------ ----------------------
node1_SSD_1 0B 23.65GB Creating
node2_SSD_1 0B 23.65GB Creating

cluster1::> storage aggregate show-auto-provision-progress
Current Target
Aggregate Usable Size Usable Size Provisioning Progress
--------------------------- ------------ ------------ ----------------------
node1_SSD_1 23.65GB 23.65GB Completed
node2_SSD_1 23.65GB 23.65GB Completed
```

```
[[ID3e74d0c33a65c1889a636e019914601a]]
= storage aggregate show-cumulated-efficiency
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display cumulated storage efficiency details

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show-cumulated-efficiency` command displays
information about the cumulated storage efficiency of all the aggregates.
The storage efficiency is displayed at four different levels:

* Total
* Aggregate
* Volume
* Snapshot and FlexClone volume



== 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.

| ``[-all-details  ]`` [.advaced-priv]#(privilege: advanced)#::
Use this parameter to show additional Storage Efficiency Ratios and size
```

```
values.

| ``[-details  ]``::
Use this parameter to show additional Storage Efficiency Ratios.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-aggregates <aggregate name>,...]`` - List of Aggregates to cumulate
Storage Efficiency ratio::
If this parameter is specified, the command calculates the cumulated
storage efficiency of the specified list of aggregates.

``[-nodes {<nodename>|local}]`` - List of Aggregates to cumulate Storage
Efficiency ratio::
If this parameter is specified, the command calculates the cumulated
storage efficiency of aggregates that are located on the specified list of
node.

``[-total-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Size Used
by volumes, clones, Snapshot copies in the Aggregate [.advaced-
priv]#(privilege: advanced)#::
Displays the total logical size used in all the specified aggregates. This
includes volumes, clones, and snapshots in all the specified aggregates.
The logical size is computed based on physical usage and savings obtained
in all the specified aggregates.

``[-total-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total Physical
Used [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by all the specified aggregates.

``[-total-storage-efficiency-ratio <text>]`` - Total Storage Efficiency
Ratio::
Displays the total storage efficiency ratio of the aggregate.

``[-total-data-reduction-logical-used-wo-snapshots
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Logical Used Without
Snapshots [.advaced-priv]#(privilege: advanced)#::
Displays the total logical size used in all the specified aggregates
excluding snapshots.

``[-total-data-reduction-physical-used-wo-snapshots
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Physical Used
Without Snapshots [.advaced-priv]#(privilege: advanced)#::
Displays the total physical size used by all the specified aggregates
```

```
aggregates excluding snapshots.

``[-total-data-reduction-efficiency-ratio-wo-snapshots <text>]`` - Total
Data Reduction Efficiency Ratio Without Snapshots::
Displays the total storage efficiency ratio obtained by Deduplication,
Compression, Data Compaction, Pattern Detection and FlexClone data
reduction technologies excluding snapshots on the specified aggregates.

``[-total-data-reduction-logical-used-wo-snapshots-flexclones
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Logical Used without
snapshots and flexclones [.advaced-priv]#(privilege: advanced)#::
Displays the total logical size used in all the specified aggregates
excluding snapshots and clones.

``[-total-data-reduction-physical-used-wo-snapshots-flexclones
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Physical Used
without snapshots and flexclones [.advaced-priv]#(privilege: advanced)#::
Displays the total physical size used by all the specified aggregates
aggregates excluding snapshots and clones.

``[-total-data-reduction-efficiency-ratio-wo-snapshots-flexclones
<text>]`` - Total Data Reduction Efficiency Ratio without snapshots and
flexclones::
Displays the total storage efficiency ratio obtained by Deduplication,
Compression, Data Compaction, Pattern Detection data reduction
technologies excluding snapshots on the specified aggregates.

``[-total-performance-tier-data-reduction-physical-used-wo-snapshots-
flexclones {<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Physical
Used without snapshots and flexclones in the FabricPool Performance Tier::
Displays the total performance tier physical size used by all the
specified aggregates aggregates excluding snapshots and clones.

``[-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Space
Used for All volumes [.advaced-priv]#(privilege: advanced)#::
Displays the total logical size used by all the volumes in all the
specified aggregates.

``[-volume-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Physical Space
Used for All volumes [.advaced-priv]#(privilege: advanced)#::
Displays the total physical size used by all volumes in all the specified
aggregates.

``[-volume-dedupe-zero-pattern-saved {<integer>[KB|MB|GB|TB|PB]}]`` -
Space Saved by volume Deduplication and pattern detection [.advaced-
priv]#(privilege: advanced)#::
```

```
Displays the total disk space that is saved by deduplication, Zero pattern
detection and FlexClone for files, LUNs or NVMe namespaces by all volumes
in all the specified aggregates.

``[-volume-efficiency-saved-ratio <text>]`` - Volume Deduplication Savings
ratio::
Displays the storage efficiency ratio for savings by deduplication and
FlexClone for files, LUNs or NVMe namespaces by all volumes in all the
specified aggregates.

``[-volume-compression-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved
by volume Compression [.advaced-priv]#(privilege: advanced)#::
Displays the total disk space that is saved by compressing blocks by all
volumes in all the specified aggregates.

``[-volume-compression-saved-ratio <text>]`` - Volume Compression Savings
ratio::
Displays the storage efficiency ratio for savings by compressing blocks on
all volumes in all the specified aggregates.

``[-volume-data-reduction-storage-efficiency-ratio <text>]`` - Volume Data
Reduction SE Ratio::
Displays the storage efficiency ratio of all the volumes in all the
specified aggregates.

``[-aggr-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Space Used
by the Aggregate [.advaced-priv]#(privilege: advanced)#::
Displays the logical size used by all the specified aggregates.

``[-aggr-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Physical Space
Used by the Aggregate [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by all the specified aggregates.

``[-aggr-data-reduction-storage-efficiency-ratio <text>]`` - Aggregate
Data Reduction SE Ratio::
Displays the storage efficiency ratio of the aggregate.

``[-snapshot-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Size
Used by Snapshot copies [.advaced-priv]#(privilege: advanced)#::
Displays the logical size used by all volume snapshots residing in all the
specified aggregates.

``[-snapshot-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Physical Size
Used by Snapshot copies [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by all volume snapshots residing in all
the specified aggregates.
```

```
``[-snapshot-volume-data-reduction-storage-efficiency-ratio <text>]`` -
Snapshot volume Data Reduction Ratio::
Displays the snapshot volume storage efficiency ratio of the aggregate.

``[-flexclone-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical
Size Used by FlexClone volumes [.advaced-priv]#(privilege: advanced)#::
Displays the logical size used by all FlexClone volumes residing in all
the specified aggregates.

``[-flexclone-volume-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` -
Physical Sized Used by FlexClone volumes [.advaced-priv]#(privilege:
advanced)#::
Displays the physical size used by all FlexClone volumes in all the
specified aggregates.

``[-flexclone-volume-data-reduction-storage-efficiency-ratio <text>]`` -
FlexClone volume Data Reduction Ratio::
Displays the FlexClone volume storage efficiency ratio of the aggregate.

``[-snapshot-flexclone-volume-data-reduction-storage-efficiency-ratio
<text>]`` - Snapshot And FlexClone volume Data Reduction SE Ratio::
Displays the snapshot and FlexClone volume storage efficiency ratio of the
aggregate.

``[-number-of-offline-volumes <integer>]`` - Number of volumes Offline::
Displays the number of volumes that are offline in all the specified
aggregates.

``[-number-of-sis-disabled-volumes <integer>]`` - Number of SIS Disabled
volumes::
Displays the number of volumes on which volume efficiency is disabled in
all the specified aggregates.

``[-number-of-sis-change-log-disabled-volumes <integer>]`` - Number of SIS
Change Log Disabled volumes [.advaced-priv]#(privilege: advanced)#::
Displays the number of volumes on which efficiency change log is disabled
in all the specified aggregates. The scheduled background Deduplication
will be disabled on these volumes.

``[-number-of-skipped-aggregates <integer>]`` - Number of Skipped
Aggregates::
Displays the number of aggregates that were skipped for calculating the
cumulated storage efficiency.

``[-skipped-aggregates <aggregate name>,...]`` - List of Aggregates
skipped::
Displays the list of aggregates that were skipped for calculating the
```

```
cumulated storage efficiency.

== Examples

The following example displays information about all aggregates that are
owned by nodes in the local cluster:
```

cluster::> aggr show-cumulated-efficiency
Total Data Reduction Efficiency Ratio: 5.00:1
Total Storage Efficiency Ratio: 6.97:1

cluster::> aggr show-cumulated-efficiency -details
Total Data Reduction Ratio: 8.44:1
Total Storage Efficiency Ratio: 6.97:1
Aggregate level Storage Efficiency
(Aggregate Deduplication and Data Compaction): 1.00:1
Volume Deduplication Efficiency: 1.12:1
Compression Efficiency: 5.73:1
Snapshot Volume Storage Efficiency: 1.00:1
FlexClone Volume Storage Efficiency: 1.00:1
Number of Offline Volumes: 0
Number of Skipped Aggregates: 0
Number of Efficiency Disabled Volumes: 0

cluster::> aggr show-cumulated-efficiency -aggregates aggr1
Total Data Reduction Efficiency Ratio: 6.00:1
Total Storage Efficiency Ratio: 7.41:1
saiscluster-1::*> aggr show-cumulated-efficiency -all-details
----- Total Data Reduction Efficiency ---------
Logical Physical Storage
Used Used Efficiency Ratio
---------- ---------- ----------
89.11MB 10.69MB 8.34:1

----- Total Storage Efficiency ----------
Logical Physical Storage
Used Used Efficiency Ratio
---------- ---------- ----------
89.11MB 12.91MB 6.90:1

— Aggregate level Storage Efficiency ---------
Logical Physical Storage
Used Used Efficiency Ratio
---------- ---------- ----------
12.91MB 12.91MB 1.00:1

-------- Volume level Storage Efficiency ----------
Logical Physical Total Volume Level Data
Used Used Reduction Efficiency Ratio
---------- ---------- ----------
84.74MB 5.51MB 15.39:1
---- Deduplication ---- ------ Compression ----

Savings Efficiency Savings Efficiency
Ratio Ratio
---------- ---------- ---------- ----------
9.27MB 1.12:1 69.96MB 5.73:1

```
 -------------Snapshot--------------------
    Logical    Physical           Storage
       Used        Used  Efficiency Ratio
----------- -----------  ----------------
         0B      2.22MB            1.00:1
 -------------FlexClone-------------------
    Logical    Physical           Storage
       Used        Used  Efficiency Ratio
----------- -----------  ----------------
         0B          0B            1.00:1
Number of Offline Volumes: 0
                       Number of Skipped Aggregates: 0
              Number of Efficiency Disabled Volumes: 0
Number of Background Deduplication Disabled Volumes: 2
```

```
[[IDf939d4ac748f46098e07db6e4227aa21]]
= storage aggregate show-efficiency
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate storage efficiency details

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show-efficiency` command displays information about
the storage efficiency of all the aggregates. The storage efficiency is
displayed at four different levels:

* Total
* Aggregate
* Volume
* Snapshot and FlexClone volume
```

```
== 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.

| ``[-advanced  ]`` [.advaced-priv]#(privilege: advanced)#::
Use this parameter to show additional Storage Efficiency Ratios and size
values.

| ``[-details  ]``::
Use this parameter to show additional Storage Efficiency Ratios.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-aggregate <aggregate name>]`` - Name of the Aggregate::
Displays the aggregate name. If this parameter is specified, the command
displays detailed information about the storage efficiency of the
specified aggregate.

``[-node {<nodename>|local}]`` - Node where Aggregate Resides::
Displays the node which owns the aggregate. If this parameter is
specified, the command displays storage efficiency information only about
the aggregates that are located on the specified node.

``[-total-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Size Used
by Volumes, Clones, Snapshot Copies in the Aggregate [.advaced-
priv]#(privilege: advanced)#::
Displays the logical size used in the aggregate. This includes volumes,
clones, and snapshots in the aggregate. The logical size is computed based
on physical usage and savings obatained in the aggregate.

``[-total-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total Physical
Used [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by the aggregate.

``[-total-storage-efficiency-ratio <text>]`` - Total Storage Efficiency
Ratio::
Displays the total storage efficiency ratio of the aggregate.

``[-total-data-reduction-logical-used-wo-snapshots
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Logical Used Without
Snapshots [.advaced-priv]#(privilege: advanced)#::
```

```
Displays the logical size used in the aggregate excluding snapshots.

``[-total-data-reduction-physical-used-wo-snapshots
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Physical Used
Without Snapshots [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by the aggregate excluding snapshots.

``[-total-data-reduction-efficiency-ratio-wo-snapshots <text>]`` - Total
Data Reduction Efficiency Ratio Without Snapshots::
Displays the total storage efficiency ratio obtained by Deduplication,
Compression, Data Compaction, Pattern Detection and FlexClone data
reduction technologies excluding snapshots on the aggregate.

``[-total-data-reduction-logical-used-wo-snapshots-flexclones
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Logical Used without
snapshots and flexclones [.advaced-priv]#(privilege: advanced)#::
Displays the logical size used in the aggregate excluding snapshots and
clones.

``[-total-data-reduction-physical-used-wo-snapshots-flexclones
{<integer>[KB|MB|GB|TB|PB]}]`` - Total Data Reduction Physical Used
without snapshots and flexclones [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by the aggregate excluding snapshots and
clones.

``[-total-data-reduction-efficiency-ratio-wo-snapshots-flexclones
<text>]`` - Total Data Reduction Efficiency Ratio without snapshots and
flexclones::
Displays the total storage efficiency ratio obtained by Deduplication,
Compression, Data Compaction, Pattern Detection data reduction
technologies excluding snapshots and FlexClone volumes on the aggregate.

``[-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Space
Used for All Volumes::
Displays the logical size used by all the volumes in the aggregate.

``[-volume-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Physical Space
Used for All Volumes::
Displays the physical size used by all volumes in the aggregate.

``[-volume-efficiency-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved
by Volume Deduplication [.advaced-priv]#(privilege: advanced)#::
Displays the total disk space that is saved by deduplication and FlexClone
for files, LUNs or NVMe namespaces by all volumes in the aggregate.

``[-volume-dedupe-zero-pattern-saved {<integer>[KB|MB|GB|TB|PB]}]`` -
```

```
Space Saved by Volume Deduplication and pattern detection [.advaced-
priv]#(privilege: advanced)#::
Displays the total disk space that is saved by deduplication, Zero pattern
detection and FlexClone for files, LUNs or NVMe namespaces by all volumes
in the aggregate.

``[-volume-efficiency-saved-ratio <text>]`` - Volume Deduplication Savings
ratio::
Displays the storage efficiency ratio for savings by deduplication and
FlexClone for files, LUNs or NVMe namespaces by all volumes in the
aggregate.

``[-volume-compression-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved
by Volume Compression [.advaced-priv]#(privilege: advanced)#::
Displays the total disk space that is saved by compressing blocks by all
volumes in the aggregate.

``[-volume-compression-saved-ratio <text>]`` - Volume Compression Savings
ratio::
Displays the storage efficiency ratio for savings by compressing blocks on
all volumes in the aggregate.

``[-volume-vbn-zero-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved by
Inline Zero Pattern Detection::
Displays the total disk space that is saved by inline zero pattern
detection by all the volumes in the aggregate.

``[-volume-data-reduction-storage-efficiency-ratio <text>]`` - Volume Data
Reduction SE Ratio::
Displays the storage efficiency ratio of all the volumes in the aggregate.

``[-aggr-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Space Used
by the Aggregate [.advaced-priv]#(privilege: advanced)#::
Displays the logical size used by the aggregate.

``[-aggr-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Physical Space
Used by the Aggregate [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by the aggregate.

``[-aggr-efficiency-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved by
Aggregate Data Reduction [.advaced-priv]#(privilege: advanced)#::
Displays the total disk space that is saved by aggregate data reduction.

``[-aggr-data-reduction-storage-efficiency-ratio <text>]`` - Aggregate
Data Reduction SE Ratio [.advaced-priv]#(privilege: advanced)#::
Displays the storage efficiency ratio of the aggregate.
```

```
``[-snapshot-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Size
Used by Snapshot Copies [.advaced-priv]#(privilege: advanced)#::
Displays the logical size used by all volume snapshots residing in the
aggregate.

``[-snapshot-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Physical Size
Used by Snapshot Copies [.advaced-priv]#(privilege: advanced)#::
Displays the physical size used by all volume snapshots residing in the
aggregate.

``[-snapshot-volume-data-reduction-storage-efficiency-ratio <text>]`` -
Snapshot Volume Data Reduction Ratio::
Displays the snapshot volume storage efficiency ratio of the aggregate.

``[-flexclone-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical
Size Used by FlexClone Volumes [.advaced-priv]#(privilege: advanced)#::
Displays the logical size used by all FlexClone volumes residing in the
aggregate.

``[-flexclone-volume-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` -
Physical Sized Used by FlexClone Volumes [.advaced-priv]#(privilege:
advanced)#::
Displays the physical size used by all FlexClone volumes in the aggregate.

``[-flexclone-volume-data-reduction-storage-efficiency-ratio <text>]`` -
FlexClone Volume Data Reduction Ratio::
Displays the FlexClone volume storage efficiency ratio of the aggregate.

``[-snapshot-flexclone-volume-data-reduction-storage-efficiency-ratio
<text>]`` - Snapshot And FlexClone Volume Data Reduction SE Ratio::
Displays the snapshot and FlexClone volume storage efficiency ratio of the
aggregate.

``[-number-of-offline-volumes <integer>]`` - Number of Volumes Offline::
Displays the number of volumes that are offline in the aggregate.

``[-number-of-sis-disabled-volumes <integer>]`` - Number of SIS Disabled
Volumes::
Displays the number of volumes on which volume efficiency is disabled in
the aggregate.

``[-number-of-sis-change-log-disabled-volumes <integer>]`` - Number of SIS
Change Log Disabled Volumes [.advaced-priv]#(privilege: advanced)#::
Displays the number of volumes on which efficiency change log is disabled
in the aggregate. The scheduled background Deduplication will be disabled
on these volumes.
```

```
== Examples

The following example displays information about all aggregates that are
owned by nodes in the local cluster:
```

cluster::*> aggr show-efficiency

Aggregate: aggr1
Node: node1

Total Data Reduction Efficiency Ratio: 3.29:1
Total Storage Efficiency Ratio: 4.29:1
Aggregate: aggr2
Node: node1

Total Data Reduction Efficiency Ratio: 4.50:1
Total Storage Efficiency Ratio: 5.49:1

cluster::*> aggr show-efficiency -details

Aggregate: aggr1
Node: node1

Total Data Reduction Ratio: 2.39:1
Total Storage Efficiency Ratio: 4.29:1

Aggregate level Storage Efficiency
(Aggregate Deduplication and Data Compaction): 1.00:1
Volume Deduplication Efficiency: 5.03:1
Compression Efficiency: 1.00:1

Snapshot Volume Storage Efficiency: 8.81:1
FlexClone Volume Storage Efficiency: 1.00:1
Number of Efficiency Disabled Volumes: 1

Aggregate: aggr2
Node: node1
Total Data Reduction Ratio: 2.39:1
Total Storage Efficiency Ratio: 4.29:1

Aggregate level Storage Efficiency
(Aggregate Deduplication and Data Compaction): 1.00:1
Volume Deduplication Efficiency: 5.03:1
Compression Efficiency: 1.00:1

Snapshot Volume Storage Efficiency: 8.81:1
FlexClone Volume Storage Efficiency: 1.00:1
Number of Efficiency Disabled Volumes: 1

```
[[ID348812cf352c1f3e4eae328997b36f08]]
= storage aggregate show-resync-status
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate resynchronization status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show-resync-status` command displays resync status
information for each plex. The command output depends on the parameter or
parameters specified with the command. If no parameters are specified, the
command displays the following information about all aggregates:

* Aggregate Name
* Resyncing Plex Name
* Resyncing Percentage



== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
This parameter specifies the name of the aggregate.

``[-plex <text>]`` - Plex Name::
This parameter specifies the name of the plex.

``[-status <text>]`` - Status::
Displays plex status. Possible values are:
+
  * ``_normal_``
  * ``_failed_``
  * ``_empty_``
```

```
  * ``_invalid_``
  * ``_uninitialized_``
  * ``_failed assimilation_``
  * ``_limbo_``
  * ``_active_``
  * ``_inactive_``
  * ``_resyncing_``

+
These values may appear by themselves or in combination separated by
commas; for example, "``_normal_`` ,``_active_`` ".

``[-is-online {true|false}]`` - Is Online::
Indicates whether the plex is online.

``[-in-progress {true|false}]`` - Resync is in Progress::
Indicates whether the plex is currently resyncing.

``[-resyncing-percent <percent>]`` - Resyncing Percentage::
Displays the resynchronization completion percentage if the plex is
currently being resynced, '-' otherwise.

``[-resync-level <integer>]`` - Resync Level::
Displays the resync level if the plex is currently being resynced, '-'
otherwise.

``[-pool <integer>]`` - Pool::
The pool number to which the majority of disks in the plex belong.

== Examples

The following example displays resynchronization status for all the
aggregates:
```

cluster1::> storage aggregate show-resync-status
Complete
Aggregate Resyncing Plex Percentage
--------- ---------- ----------
aggr0 plex0 -
aggr1 plex0 -
aggr1 plex1 10.00
aggr2 plex0 -
aggr2 plex2 -
5 entries were displayed.

```
[[ID77d81eb21f30d9eb3b2ee3d39b8450af]]
= storage aggregate show-scrub-status
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate scrubbing status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show-scrub-status` command displays the following
information about the scrub status of aggregates:

* Aggregate name
* RAID groups
* Whether the scrub is suspended
* Percentage of the scrub that is completed
* Last scrub time of the aggregate



== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
If this parameter is specified, the command displays detailed scrub-status
information about the specified aggregate.

``[-raidgroup <text>]`` - RAID Group::
If this parameter is specified, the command displays information only
about the aggregate that contains the specified RAID group.

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays information only
about the aggregates on the specified node. The value ``_local_``
specifies the current node.
```

```
``[-suspended {true|false}]`` - Is Suspended::
If this parameter is specified, the command displays information only
about the aggregates that have the specified scrub-suspension state (true
or false).

``[-complete-percentage <percent>]`` - Percentage Completed::
If this parameter is specified, the command displays information only
about the aggregates whose scrubs have the specified completed percentage.

``[-last-scrub-time <MM/DD/YYYY HH:MM:SS>]`` - Last Scrub Time::
If this parameter is specified, the command displays information only
about the aggregates that have the specified last-scrub time, in the
format ``_MM/DD/YYYY HH:MM:SS_`` .

== Examples

The following example displays scrub-status information for all the
aggregates:
```

cluster1::> storage aggregate show-scrub-status
Aggregate RAID Groups Suspended Percentage Last Scrub Time
--------- ---------- --------- --------- ----------
aggr0 /aggr0/plex0/rg0 true 0% 3/31/2011 21:23:02
aggr1 /aggr1/plex0/rg1 true 45% 3/30/2011 01:05:00
aggr2 /aggr2/plex0/rg0 true 33% 3/30/2011 23:43:34
aggr3 /aggr3/plex0/rg1 true 79% 3/29/2011 00:34:36
4 entries were displayed.

```
The following example displays detailed information about the aggregate
named aggr1:
```

cluster1::> storage aggregate show-scrub-status -instance -aggregate aggr1
Aggregate: aggr1
RAID Group: /aggr1/plex0/rg0
Is Suspended: false
Percentage Completed: 2%
Last Scrub Time: 3/31/2011 22:02:50

```
[[ID36bf1baba605db6c152c3a5f41daff13]]
= storage aggregate show-space
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display details of space utilization within an aggregate.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show-space` command displays information about
space utilization within aggregates and any attached external capacity
tier. The command output breaks down space usage in the specified
aggregate by feature. If no parameters are specified, the command displays
this information about all aggregates. Note that used percentage for an
external capacity tier will be non-zero only if a size limit was set for
that aggregate's attached tier.

== 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.

``[-aggregate-name <aggregate name>]`` - Aggregate::
If this parameter is specified, the command displays information only
about the specified aggregates.

``[-bin-num <integer>]`` - Bin Number::
If this parameter is specified, the command displays information only
about the aggregates whose bin number for the storage tier matches the
specified value. Typically, bin 0 refers to the performance tier or active
file system and bin numbers greater than 0 refer to the external capacity
tiers attached to the aggregate.

``[-tier-name <text>]`` - Tier Name For Show Command::
If this parameter is specified, the command displays information only
about the aggregates whose attached storage tier name matches the
specified value.

``[-aggregate <aggregate name>]`` - Aggregate Display Name::
If this parameter is specified, the command displays information only
about space used in the specified aggregate or aggregates.
```

```
``[-aggregate-uuid <UUID>]`` - Uuid of the Aggregate::
If this parameter is specified, the command displays information only
about the aggregates whose UUID matches the specified value.

``[-volume-footprints {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Footprints::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified amount of space
in use by volume footprints. A volume's footprint is the overall amount of
space that a volume occupies in the aggregate, including the volume
metadata and data.

``[-volume-footprints-percent <percent_no_limit>]`` - Volume Footprints
Percent::
If this parameter is specified, the command displays information only
about the aggregate or aggregates whose volume footprints occupy the
specified percentage of space.

``[-snap-size-total {<integer>[KB|MB|GB|TB|PB]}]`` - Total Space for
Snapshot Copies in Bytes::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified amount of space
in use by aggregate snapshots. This field includes the space that is
reserved for snapshots and is not available to volumes or aggregate data
and metadata. It is set to 0 by default.

``[-percent-snapshot-space <percent>]`` - Space Reserved for Snapshot
Copies::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified percentage of
space in use by aggregate snapshots.

``[-aggregate-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - Aggregate
Metadata::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified amount of space
in use by aggregate metadata.

``[-aggregate-metadata-percent <percent_no_limit>]`` - Aggregate Metadata
Percent::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified percentage of
space in use by aggregate metadata.

``[-used-including-snapshot-reserve {<integer>[KB|MB|GB|TB|PB]}]`` - Total
Used::
If this parameter is specified, the command displays information only
```

```
about the aggregate or aggregates that have the specified amount of space
in use in the aggregate.
+
It is important to note that this parameter treats the entire snapshot
reserve as used space since it is not available for volumes.

``[-used-including-snapshot-reserve-percent <percent_no_limit>]`` - Total
Used Percent::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified percentage of
space in use in the aggregate and its snapshot reserve.

``[-aggregate-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified size.

``[-snapshot-reserve-unusable {<integer>[KB|MB|GB|TB|PB]}]`` - Snapshot
Reserve Unusable::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified amount of space
reserved but unusable in the volume.
+
snapshot reserve can be diminished under certain conditions to accommodate
volume metadata. Creating space in the aggregate will make this space
available.

``[-snapshot-reserve-unusable-percent <percent_no_limit>]`` - Snapshot
Reserve Unusable Percent::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified percentage of
space reserved but unusable.

``[-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total Physical Used
Size::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified amount of
physical space in use by the aggregate.
+
This differs from ``_total-used_`` space by the space that is guaranteed
for future writes. The value includes blocks in use by snapshots.

``[-physical-used-percent <percent_no_limit>]`` - Physical Used
Percentage::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified percentage of
physical space in use in the aggregates.
```

```
``[-performance-tier-inactive-user-data {<integer>[KB|MB|GB|TB|PB]}]`` -
Performance Tier Inactive User Data::
If this parameter is specified, the command displays information only
about the aggregates whose amount of inactive user data in the performance
tier matches the specified value. The inactive user data can be tiered out
to a capacity tier if the aggregate is a FabricPool.

``[-performance-tier-inactive-user-data-percent <percent>]`` - Performance
Tier Inactive User Data Percent::
If this parameter is specified, the command displays information only
about the aggregates whose percentage of inactive user data in the
performance tier matches the specified value.

``[-cross-volume-dedupe-metadata {<integer>[KB|MB|GB|TB|PB]}]`` -
Aggregate Dedupe Metadata::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified amount of space
in use by cross volume deduplication metadata.

``[-cross-volume-dedupe-metadata-percent <percent_no_limit>]`` - Aggregate
Dedupe Metadata Percent::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified percentage of
space in use by cross volume deduplication metadata.

``[-cross-volume-dedupe-temp-metadata {<integer>[KB|MB|GB|TB|PB]}]`` -
Aggregate Dedupe Temporary Metadata::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified amount of space
in use by cross volume deduplication temporary metadata.

``[-cross-volume-dedupe-temp-metadata-percent <percent_no_limit>]`` -
Aggregate Dedupe Temporary Metadata Percent::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that have the specified percentage of
space in use by cross volume deduplication temporary metadata.

``[-total-provisioned-space {<integer>[KB|MB|GB|TB|PB]}]`` - Total Space
Provisioned inside Aggregate::
If this parameter is specified, the command displays information about the
total provisioned space across all the volumes in the aggr.

``[-total-provisioned-space-percent <percent_no_limit>]`` - Percentage
Space Provisioned inside Aggregate::
If this parameter is specified, the command displays information about the
percentage of total provisioned space across all the volumes against
```

```
aggr's size.

``[-object-store-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total
Object Store Physical Used Size::
If this parameter is specified, the command displays information only
about the aggregates whose physical space use in the attached object store
matches the specified value.

``[-object-store-physical-used-percent <percent_no_limit>]`` - Object
Store Physical Used Percentage::
If this parameter is specified, the command displays information only
about aggregates whose physical space in use in the attached object store
as a percentage of the license limit matches the specified value.

``[-object-store-referenced-capacity {<integer>[KB|MB|GB|TB|PB]}]`` -
Total Object Store Logical Referenced Capacity::
If this parameter is specified, the command displays information only
about the aggregates whose logical reference capacity space in use in the
attached object store matches the specified value.

``[-object-store-referenced-capacity-percent <percent_no_limit>]`` -
Object Store Logical Referenced Capacity Percentage::
If this parameter is specified, the command displays information only
about aggregates whose logical reference capacity space in use in the
attached object store as a percentage of the license limit matches the
specified value.

``[-object-store-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - (DEPRECATED)-
Object Store Metadata::
This parameter is deprecated in ONTAP 9.10.1 and later. If this parameter
is specified, the command displays information only about the aggregates
whose metadata space in use in the attached object store matches the
specified value.

``[-object-store-metadata-percent <percent_no_limit>]`` - (DEPRECATED)-
Object Store Metadata Percent::
This parameter is deprecated in ONTAP 9.10.1 and later. If this parameter
is specified, the command displays information only about aggregates whose
metadata space in use in the attached object store as a percentage of the
license limit matches the specified value.

``[-object-store-unreclaimed-space {<integer>[KB|MB|GB|TB|PB]}]`` -
(DEPRECATED)-Total Unreclaimed Space::
This parameter is deprecated in ONTAP 9.10.1 and later. If this parameter
is specified, the command displays information only about the aggregates
whose unreclaimed space in use in the attached object store matches the
```

```
specified value.

``[-object-store-unreclaimed-space-percent <percent_no_limit>]`` -
(DEPRECATED)-Object Store Unreclaimed Space Percentage::
This parameter is deprecated in ONTAP 9.10.1 and later. If this parameter
is specified, the command displays information only about aggregates whose
unreclaimed space in use in the attached object store as a percentage of
the license limit matches the specified value.

``[-object-store-size {<integer>[KB|MB|GB|TB|PB]}]`` - Object Store Size::
If this parameter is specified, the command displays information only
about the aggregates whose attached object store size limit matches the
specified value.

``[-object-store-sis-space-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Object
Store Space Saved by Storage Efficiency::
If this parameter is specified, the command displays information only
about the aggregates whose amount of space saved by storage efficiency
matches the specified value.

``[-object-store-sis-space-saved-percent <percent_no_limit>]`` - Object
Store Space Saved by Storage Efficiency Percentage::
If this parameter is specified, the command displays information only
about the aggregates whose percentage of space saved by storage efficiency
matches the specified value.

``[-object-store-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total
Object Store Logical Used Size::
If this parameter is specified, the command displays information only
about the aggregates whose logical space use in the attached object store
matches the specified value.

``[-object-store-logical-used-percent <percent_no_limit>]`` - Object Store
Logical Used Percentage::
If this parameter is specified, the command displays information only
about aggregates whose logical space in use in the attached object store
as a percentage of the license limit matches the specified value.

``[-object-store-logical-unreferenced-capacity
{<integer>[KB|MB|GB|TB|PB]}]`` - Object Store Logical Unreferenced
Capacity::
If this parameter is specified, the command displays information only
about the aggregates whose logical unreferenced capacity in use in the
attached object store matches the specified value.

``[-object-store-logical-unreferenced-capacity-percent
```

```
<percent_no_limit>]`` - Object Store Logical Unreferenced Percentage::
This parameter is deprecated in ONTAP 9.10.1 and later. If this parameter
is specified, the command displays information only about aggregates whose
logical unreferenced capacity in use in the attached object store as a
percentage of the license limit matches the specified value.

== Examples

The following example displays information about all aggregates:
```

cluster1::> storage aggregate show-space
Aggregate : aggr0
Feature Used Used%
--- --- ---
Volume Footprints 5.75GB 91%
Aggregate Metadata 380KB 0%
Snapshot Reserve 325.3MB 5%
Total Used 6.07GB 96%
Total Physical Used 221.9MB 3%
Aggregate : aggr1
Feature Used Used%
--- --- ---
Volume Footprints 2.03GB 33%
Aggregate Metadata 304KB 0%
Total Used 2.03GB 33%
Total Physical Used 2.23MB 0%

2 entries were displayed.

```
The following example displays information about all the aggregates in a
system including the ones that have an object store attached to them.
```

cluster-1::> storage aggregate show-space
Aggregate : aggr0
Feature Used Used%
--- --- ---
Volume Footprints 2.87GB 90%
Aggregate Metadata 328KB 0%
Snapshot Reserve 162.6MB 5%
Total Used 3.03GB 95%
Total Physical Used 2.08GB 65%
Aggregate : aggr1
Performance Tier
Feature Used Used%
--- --- ---
Volume Footprints 1.25GB 13%
Aggregate Metadata 540KB 0%
Snapshot Reserve 0B 0%
Total Used 1.25GB 13%
Total Physical Used 1.23GB 13%

Aggregate : aggr1
Object Store: my-store
Feature Used Used%
--- --- ---
Referenced Capacity 811.2MB 0%
Metadata 0B 0%
Unreclaimed Space 0B 0%
Space Saved by Storage Efficiency 0B 0%
Total Physical Used 811.2MB 0%

```
[[IDea66b7ffab022b1c7e0c8e43aa4d56f5]]
= storage aggregate show-spare-disks
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display spare disks

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The command `storage aggregate show-spare-disks` displays information
about spare disks. The command output depends on the parameter or
parameters specified with the command. If no parameters are specified, the
command displays information about all spare disks in the cluster.

== 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.

| ``[-partition-info  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information about root-data and root-data1-data2
partitioned spares.
+
  * Disk
  * Type
  * Class
  * RPM
  * Checksum
  * Local Data Usable
```

```
  * Local Data1 Usable
  * Local Data2 Usable
  * Local Root Usable
  * Physical Size
  * Status



| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed information
about each spare disk.

``[-original-owner <text>]`` - Original Owner::
Selects the spare disks that match this parameter value.

``[-disk <disk path name>]`` - Disk Name::
Selects the spare disks that match this parameter value.

``[-checksum-style {advanced_zoned | block | none}]`` - Checksum Style::
Selects the spare disks that match this parameter value. Possible values
are:
+
  * block -- Supports block checksum
  * advanced_zoned -- Supports advanced zone checksum
  * none -- No checksum support



``[-disk-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS  | SSD |
VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Disk Type::
Selects the spare disks that match this parameter value.

``[-effective-disk-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS  |
SSD | VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` -
Effective Disk Type::
Selects the spare disks that match this parameter value.
+
Hard disk drives with the same `effective-disk-type` value may be mixed
together in the same aggregate depending upon the system's
``_raid.mix.hdd.disktype.capacity_`` and
``_raid.mix.hdd.disktype.performance_`` option settings. Solid state
drives with the same `effective-disk-type` value may be mixed together in
the same aggregate depending upon the system's
``_raid.mix.disktype.solid_state_`` option setting.

``[-standard-disk-type {SATA | FC | NL-SAS | LUN | SAS | SCSI | SSD | VM-
DISK | NVMe-SSD | SSD-CAP | SSD-ZNS | VM-LUN | VM-LUN-SSDS}]`` - Standard
```

```
Disk Type::
Selects the spare disks that match this parameter value.

``[-disk-class {capacity | performance | archive | solid-state | array |
virtual}]`` - Disk Class::
Selects the spare disks that match this parameter value. Possible values
are:
+
  * capacity -- Capacity-oriented, near-line disk types. Includes disk
types FSAS, BSAS and ATA.
  * performance -- Performance-oriented, enterprise class disk types.
Includes disk types FCAL and SAS.
  * archive -- Archive class SATA disks in multi-disk carrier storage
shelves. Includes disk type MSATA.
  * solid-state -- Solid-state drives. Includes disk type SSD, SSD-CAP and
SSD-NVM.
  * array -- Logical storage devices backed by storage arrays and used by
ONTAP as disks. Includes disk type LUN.
  * virtual -- Virtual disks that are formatted and managed by the
hypervisor. Includes disk type VMDISK.



+
Disks with the same `disk-class` value are compatible for use in the same
aggregate.

``[-disk-rpm <integer>]`` - Disk RPM::
Selects the spare disks that match this parameter value.

``[-effective-disk-rpm <integer>]`` - Effective Disk RPM::
Selects the spare disks that match this parameter value.
+
Hard disk drives with the same `effective-disk-rpm` value may be mixed
together in the same aggregate depending upon the system's
``_raid.mix.hdd.rpm.capacity_`` and ``_raid.mix.hdd.rpm.performance_``
option settings.

``[-syncmirror-pool <text>]`` - Pool Number::
Selects the spare disks that match this parameter value.

``[-owner-name {<nodename>|local}]`` - Current Owner::
Selects the spare disks that match this parameter value.

``[-home-owner-name {<nodename>|local}]`` - Home Owner::
Selects the spare disks that match this parameter value.
```

```
``[-dr-owner-name {<nodename>|local}]`` - DR Home Owner::
Selects the spare disks that match this parameter value.

``[-usable-size-blks <integer>]`` - Disk Usable Size in 4K blocks::
Selects the spare disks that match this parameter value.

``[-local-usable-data-size-blks <integer>]`` - Local Node Data Usable Size
in 4K blocks::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
one or two data aggregates.
+
This value describes the data partition size (of root-data partitioned
disk) or the combined data1 + data2 partition size (of root-data1-data2
partitioned disk) in 4KB blocks.

``[-local-usable-root-size-blks <integer>]`` - Local Node Root Usable Size
in 4K blocks::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
one or two data aggregates.
+
This value describes the root partition size in 4KB blocks.

``[-usable-size {<integer>[KB|MB|GB|TB|PB]}]`` - Disk Usable Size::
Selects the spare disks that match this parameter value.

``[-total-size {<integer>[KB|MB|GB|TB|PB]}]`` - Total Size::
Selects the spare disks that match this parameter value.

``[-local-usable-data-size {<integer>[KB|MB|GB|TB|PB]}]`` - Local Node
Data Usable Size::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
```

```
one or two data aggregates.
+
This value describes the data partition size (of root-data partitioned
disk) or the combined data1 + data2 partition size (of root-data1-data2
partitioned disk) in auto-scaled units.

``[-local-usable-root-size {<integer>[KB|MB|GB|TB|PB]}]`` - Local Node
Root Usable Size::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
one or two data aggregates.
+
This value describes the root partition size in auto-scaled units.

``[-is-disk-zeroed {true|false}]`` - Is Disk Zeroed?::
Selects the spare disks that match this parameter value.
+
When disks are zeroed, they can be provisioned directly into aggregates
which avoids a lengthy zeroing process.

``[-is-disk-zeroing {true|false}]`` - Is Disk Zeroing?::
Selects the spare disks that match this parameter value.

``[-zeroing-percent <percent>]`` - Zeroing Percentage Completed::
Selects the spare disks that match this parameter value.

``[-is-sparecore {true|false}]`` - Sparecore Disk?::
Selects the spare disks that match this parameter value.

``[-sparecore-status <Spare core status>]`` - Sparecore Status::
Selects the spare disks that match this parameter value.

``[-sparecore-percent <percent>]`` - Sparecore Percentage Completed::
Selects the spare disks that match this parameter value.

``[-is-disk-shared {true|false}]`` - Is Disk Shared?::
Selects the spare disks that match this parameter value.
+
Shared disks have partitions that allow them to be used in multiple
aggregates and between nodes in an HA pair. When set to ``_true_`` , this
parameter selects shared disks in which the root partition and/or the data
partition is a spare. When set to ``_false_`` only spare disks without
```

```
partitions are displayed. When this parameter is not used, all spare disks
are displayed.

``[-is-disk-offline {true|false}]`` - Is Disk Offline?::
Selects the spare disks that match this parameter value.
+
Disk offline events are typically temporary events which allow ONTAP to
perform background error recovery activity.

``[-is-disk-sick {true|false}]`` - Is Disk Sick?::
Selects the spare disks that match this parameter value.
+
A sick disk triggers Rapid RAID Recovery to copy data to a spare drive. At
the end of the process the sick disk is marked as ``_broken_`` .

``[-is-disk-left-behind {true|false}]`` - Is Disk Left Behind Spare?::
Selects the spare disks that match this parameter value.
+
Disks are left behind if they are not responding during a giveback or
switchback event.

``[-local-usable-data1-size-blks <integer>]`` - Local Node Data1 Usable
Size in 4K blocks [.advaced-priv]#(privilege: advanced)#::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
one or two data aggregates.
+
This value describes the data1 partition size of a root-data1-data2
partitioned disk in 4KB blocks.

``[-local-usable-data2-size-blks <integer>]`` - Local Node Data2 Usable
Size in 4K blocks [.advaced-priv]#(privilege: advanced)#::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
one or two data aggregates.
+
This value describes the data2 partition size of a root-data1-data2
partitioned disk in 4KB blocks.
```

```
``[-local-usable-data1-size {<integer>[KB|MB|GB|TB|PB]}]`` - Local Node
Data1 Usable Size [.advaced-priv]#(privilege: advanced)#::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
one or two data aggregates.
+
This value describes the data1 partition size of a root-data1-data2
partitioned disk in auto-scaled units.

``[-local-usable-data2-size {<integer>[KB|MB|GB|TB|PB]}]`` - Local Node
Data2 Usable Size [.advaced-priv]#(privilege: advanced)#::
Selects the spare disks that match this parameter value.
+
Disks that have two partitions can be used for one root aggregate and one
data aggregate.
+
Disks that have three partitions can be used for one root aggregate and
one or two data aggregates.
+
This value describes the data2 partition size of a root-data1-data2
partitioned disk in auto-scaled units.

== Examples
```

```
Display spare disks owned by node ``_node-b_`` .
```

cluster1::> storage aggregate show-spare-disks -owner-name node-b

Original Owner: node-b
Pool0
Spare Pool
Usable Physical
Disk Type Class RPM Checksum Size Size Status
---------- ----- ---------- ------ ---------- -------- -------- --------
1.1.13 BSAS capacity 7200 block 827.7GB 828.0GB zeroed
1.1.15 BSAS capacity 7200 block 413.2GB 414.0GB zeroed

Original Owner: node-b
Pool0
Partitioned Spares
Local Local
Data Root Physical
Disk Type Class RPM Checksum Usable Usable Size Status
---------- ----- ---------- ------ ---------- -------- -------- -------- --------
1.0.8 SAS performance 10000 block 472.9GB 73.89GB 547.1GB zeroed

```
Check on the progress of a previous disk zeroing command.
```

cluster1::> storage aggregate show-spare-disks -owner-name node-b -zeroing-percent >0

Original Owner: node-b
Pool0
Spare Pool
Usable Physical
Disk Type Class RPM Checksum Size Size Status
---------- ----- ---------- ------ ---------- -------- -------- --------
1.1.13 BSAS capacity 7200 block 827.7GB 828.0GB zeroing, 17% done
1.1.15 BSAS capacity 7200 block 413.2GB 414.0GB zeroing, 28% done
2 entries were displayed.

```
[[ID908a0ed988aef558e72f7c0d2a6b5876]]
= storage aggregate show-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show-status` command displays the RAID layout and
disk configuration of aggregates. The command output depends on the
parameter or parameters specified with the command. If no parameters are
specified, the command displays information about all aggregates in the
cluster.

NOTE: This command does not use pagination. To reduce the output,
parameter-based filtering can be used.



== 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  ]`` }::
This parameter currently has no effect.

``[-aggregate <text>]`` - Aggregate Name::
Selects the aggregates that match this parameter value.

``[-node <nodename>]`` - Node::
Selects the aggregates that match this parameter value.

``[-aggregate-uuid <UUID>]`` - Aggregate UUID::
Selects the aggregates that match this parameter value.

== Examples

Display the RAID layout of a Flash Pool aggregate.
```

cluster1::> storage aggregate show-status -aggregate nodeB_flashpool_1

Owner Node: node-b
Aggregate: nodeB_flashpool_1 (online, raid_dp, hybrid) (block checksums)
Plex: /nodeB_flashpool_1/plex0 (online, normal, active, pool0)
RAID Group /nodeB_flashpool_1/plex0/rg0 (normal, block checksums)
Usable Physical
Position Disk Pool Type RPM Size Size Status
-------- ---------- ---- ----- ------ -------- -------- ----------
dparity 1.1.7 0 BSAS 7200 827.7GB 828.0GB (normal)
parity 1.1.8 0 BSAS 7200 827.7GB 828.0GB (normal)
data 1.1.10 0 BSAS 7200 827.7GB 828.0GB (normal)
data 1.1.11 0 BSAS 7200 827.7GB 828.0GB (normal)
data 1.1.12 0 BSAS 7200 827.7GB 828.0GB (normal)
RAID Group /nodeB_flashpool_1/plex0/rg1 (normal, block checksums) (Storage Pool: SP2)
Usable Physical
Position Disk Pool Type RPM Size Size Status
-------- ---------- ---- ----- ------ -------- -------- ----------
shared 1.0.22 0 SSD - 186.2GB 745.2GB (normal)
shared 1.0.20 0 SSD - 186.2GB 745.2GB (normal)
shared 1.0.18 0 SSD - 186.2GB 745.2GB (normal)
shared 1.0.16 0 SSD - 186.2GB 745.2GB (normal)

```
[[IDb35d0c7377a252e579ca0de963b01d8c]]
= storage aggregate show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of aggregates
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate show` command displays information about
aggregates. The command output depends on the parameter or parameters
specified with the command. If no parameters are specified, the command
displays the following information about all aggregates:

* Aggregate name
* Size
* Available size
* Percentage used
* State
* Number of volumes
* Node on which the aggregate is located
* RAID status

To display detailed information about a single aggregate, use the `-
aggregate` parameter.

== 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.

| ``[-checksum  ]``::
If this parameter is specified, the command displays information about the
checksum for all aggregates in the cluster:
+
  * Aggregate name
  * Checksum status (active, off, reverting, none, unknown, initializing,
reinitializing, reinitialized, upgrading_phase1, upgrading_phase2)
  * Checksum style (none, advanced_zoned, block, mixed, WAFL, or unknown)



| ``[-disk  ]``::
If this parameter is specified, the command displays disk names for all
aggregates in the cluster:
+
  * Aggregate name
  * Number and names of disks in the aggregate
```

```
| ``[-raid-info  ]``::
If this parameter is specified, the command displays information about
RAID groups, RAID type, maximum RAID size, checksum state, checksum style
and whether the RAID status is inconsistent.

| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed information
about all aggregates in the cluster.

``[-aggregate <aggregate name>]`` - Aggregate::
If this parameter is specified, the command displays detailed information
about the specified aggregate.

``[-storage-type {hdd | hybrid | lun | ssd | vmdisk}]`` - Storage Type::
If this parameter is specified, the command displays information only
about the aggregates with the specified storage type. The possible values
are hdd, hybrid, lun, ssd and vmdisk.

``[-chksumstyle <aggrChecksumStyle>]`` - Checksum Style::
If this parameter is specified, the command displays information only
about the aggregates that use the specified checksum style.

``[-diskcount <integer>]`` - Number Of Disks::
If this parameter is specified, the command displays information only
about the aggregates that have the specified number of disks.

``[-m, -mirror <true>]`` - Is Mirrored::
If this parameter is specified, the command displays information only
about the aggregates that have the specified mirrored value.

``[-d, -disklist <disk path name>,...]`` - Disks for First Plex::
If this parameter is specified, the command displays information only
about the aggregates that have the specified disk or disks.

``[-mirror-disklist <disk path name>,...]`` - Disks for Mirrored Plex::
If this parameter is specified, the command displays information only
about the aggregates that have the specified disk or disks present in the
mirrored plex.

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays information only
about the aggregates that are located on the specified node.

``[-free-space-realloc {on|off}]`` - Free Space Reallocation::
```

```
If this parameter is specified, the command displays whether free space
reallocation is enabled on the specified aggregate.

``[-ha-policy {sfo|cfo}]`` - HA Policy::
This optionally specifies the high-availability policy to be used in the
context of a root recovery procedure. Do not modify this setting unless
directed to do so by a customer support representative.

``[-percent-snapshot-space <percent>]`` - Space Reserved for Snapshot
Copies::
If this parameter is specified, the command displays information only
about the aggregates that have the specified space reserved for snapshots.

``[-space-nearly-full-threshold-percent <percent>]`` - Aggregate Nearly
Full Threshold Percent::
If this parameter is specified, the command displays information only
about the aggregates that have the specified nearly full threshold
percent.

``[-space-full-threshold-percent <percent>]`` - Aggregate Full Threshold
Percent::
If this parameter is specified, the command displays information only
about the aggregates that have the specified full threshold percent.

``[-hybrid-enabled {true|false}]`` - Hybrid Enabled::
If this parameter is specified, the command displays information only
about the aggregates that are eligible to contain both SSD and non-SSD
RAID groups.

``[-availsize {<integer>[KB|MB|GB|TB|PB]}]`` - Available Size::
If this parameter is specified, the command displays information only
about the aggregates that have the specified available size.

``[-chksumenabled {true|false}]`` - Checksum Enabled::
If this parameter is specified, the command displays information only
about the aggregates that have the specified checksum setting.

``[-chksumstatus <text>]`` - Checksum Status::
If this parameter is specified, the command displays information only
about the aggregates that have the specified checksum status. The possible
values for checksum status include the following: active, off, reverting,
none, unknown, initializing, reinitializing, reinitialized,
upgrading_phase1, and upgrading_phase2.

``[-cluster <text>]`` - Cluster::
If this parameter is specified, the command displays information only
about the aggregates that are owned by nodes in the specified cluster. By
```

```
default, only local cluster aggregates are displayed.

``[-cluster-id <UUID>]`` - Home Cluster ID::
If this parameter is specified, the command displays information only
about the aggregates that are owned by nodes in the cluster specified by
the cluster UUID. By default, only local cluster aggregates are displayed.

``[-dr-home-id <integer>]`` - DR Home ID::
If this parameter is specified, the command displays information only
about the aggregates whose Disaster Recovery home node has the specified
system ID.

``[-dr-home-name <text>]`` - DR Home Name::
If this parameter is specified, the command displays information only
about the aggregates whose Disaster Recovery home is the specified node.

``[-inofile-version <integer>]`` - Inofile Version [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregates whose inode files are at the specified version.

``[-has-mroot {true|false}]`` - Has Mroot Volume::
If this parameter is specified, the command displays information about
only the aggregates that contain their owning node's management root
directory.

``[-has-partner-mroot {true|false}]`` - Has Partner Node Mroot Volume::
If this parameter is specified, the command displays information about
only the aggregates that contain the management root directory of their
owning node's failover partner.

``[-home-id <integer>]`` - Home ID::
If this parameter is specified, the command displays information only
about the aggregates whose home node has the specified system ID.

``[-home-name <text>]`` - Home Name::
If this parameter is specified, the command displays information only
about the aggregates whose home node is the specified node.

``[-hybrid-cache-size-total {<integer>[KB|MB|GB|TB|PB]}]`` - Total Hybrid
Cache Size::
If this parameter is specified, the command displays information only
about the aggregates that have the specified total cache size in a Flash
Pool.

``[-hybrid {true|false}]`` - Hybrid::
If this parameter is specified, the command displays information only
```

```
about the aggregates that currently contain both SSD and non-SSD RAID
groups.

``[-inconsistent {true|false}]`` - Inconsistent::
If this parameter is specified, the command displays information only
about the aggregates that have the specified consistency.

``[-is-home {true|false}]`` - Is Aggregate Home::
If this parameter is specified, the command displays information only
about the aggregates whose home node and owner node have the same system
ID.

``[-s, -maxraidsize <integer>]`` - Max RAID Size::
If this parameter is specified, the command displays information only
about the aggregates that have the specified maximum number of disks for
RAID groups.
+
NOTE: For Flash Pool aggregates, this option controls the maximum size of
the HDD RAID groups.
+


``[-cache-raid-group-size <integer>]`` - Flash Pool SSD Tier Maximum RAID
Group Size::
If this parameter is specified, the command displays information about the
maximum RAID group size for the SSD tier for Flash Pool aggregates.
+
NOTE: This parameter is applicable only for Flash Pool aggregates.
+


``[-owner-id <integer>]`` - Owner ID::
If this parameter is specified, the command displays information only
about the aggregates that are owned by the node with the specified system
ID.

``[-owner-name <text>]`` - Owner Name::
If this parameter is specified, the command displays information only
about the aggregates that are owned by the specified node.

``[-percent-used <percent>]`` - Used Percentage::
If this parameter is specified, the command displays information only
about the aggregates that have the specified used size, as a percentage.

``[-plexes <text>,...]`` - Plexes::
If this parameter is specified, the command displays information only
about the aggregates that have the specified plex or plexes.
```

```
``[-raidgroups <text>,...]`` - RAID Groups::
If this parameter is specified, the command displays information only
about the aggregates that have the specified RAID group or groups.

``[-raidstatus <text>]`` - RAID Status::
If this parameter is specified, the command displays information only
about the aggregates that have the specified RAID status. The possible
values for RAID status are normal, copying, ironing, degraded, mirror
degraded, growing, initializing, invalid, needs check, partial,
reconstruct, raid4, raid0, raid_dp, raid_tec, redirect, and wafl
inconsistent. You can specify multiple values (for example, reconstruct
and growing).

``[-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}]`` - RAID Type::
If this parameter is specified, the command displays information only
about the aggregates that use the specified RAID type. The possible values
are ``_raid0_`` for RAID 0, ``_raid4_`` for RAID4, ``_raid_dp_`` for RAID-
DP, ``_raid_tec_`` for RAID-TEC, and ``_mixed_raid_type_`` for aggregates
that include a mix of RAID types.

``[-resyncsnaptime <integer>]`` - SyncMirror Resync Snapshot Frequency in
Minutes::
If this parameter is specified, the command displays information only
about the aggregates whose SyncMirror Resynchronization Snapshot Frequency
is the specified value.

``[-root {true|false}]`` - Is Root::
If this parameter is specified, the command displays information about
only the root aggregates in the cluster.

``[-sis-metadata-space-used {<integer>[KB|MB|GB|TB|PB]}]`` - Space Used by
Metadata for Volume Efficiency::
If this parameter is specified, the command displays information about
only the aggregates with the specified space used by A-SIS metafiles for
volume efficiency. This parameter is deprecated in ONTAP 8.2 and later.
Use the volume-footprint-list-info API for details related to space usage
by deduplication metadata

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size::
If this parameter is specified, the command displays information only
about the aggregates that have the specified size. The size of the
aggregate is reported as the size available for use by WAFL, excluding
WAFL reserve and aggregate snapshot reserve capacity. Use the
xref:{relative_path}storage-aggregate-show-space.html[storage aggregate
show-space] command to see the details of space utilization within an
```

```
aggregate.

``[-state <aggregate state>]`` - State::
If this parameter is specified, the command displays information only
about the aggregates that have the specified state.

``[-usedsize {<integer>[KB|MB|GB|TB|PB]}]`` - Used Size::
If this parameter is specified, the command displays information only
about the aggregates that have the specified used size.

``[-uses-shared-disks {true|false}]`` - Uses Shared Disks::
Selects the aggregates that match this parameter value. This parameter is
used to list all the aggregates that use shared HDDs or shared SSDs.

``[-uuid <text>]`` - UUID String [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregate that has the specified UUID. This parameter is
available only at the advanced privilege level and higher.

``[-volcount <integer>]`` - Number Of Volumes::
If this parameter is specified, the command displays information only
about the aggregates that have the specified number of volumes.

``[-is-autobalance-eligible {true|false}]`` - Is Eligible for Auto Balance
Aggregate [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregates that are considered by the Auto Balance Aggregate
feature.

``[-autobalance-state <Auto Balance Aggregate state>]`` - State of the
aggregate being balanced [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregates that have the specified state.

``[-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total Physical Used
Size::
If this parameter is specified, the command displays information only
about the aggregates that have the specified physical used size. This
differs from ``_total-used_`` space by the space that is guaranteed for
future writes. The value includes blocks in use by snapshots.

``[-physical-used-percent <percent_no_limit>]`` - Physical Used
Percentage::
If this parameter is specified, the command displays information only
about the aggregates that have the specified physical used percent.

``[-autobalance-state-change-counter <integer>]`` - State Change Counter
```

```
for Auto Balancer [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregates that have the specified number of state change caused
by the Auto Balance Aggregate feature.

``[-L, -snaplock-type {non-snaplock|compliance|enterprise}]`` - SnapLock
Type::
If this parameter is specified, the command displays information only
about the aggregates that have the specified snaplock-type.

``[-is-nve-capable {true|false}]`` - Is NVE Capable::
This parameter indicates whether or not the aggregate is capable of
supporting NVE (NetApp volume encryption).

``[-is-cft-precommit {true|false}]`` - Is in the precommit phase of Copy-
Free Transition [.advaced-priv]#(privilege: advanced)#::
Selects the aggregates that are set with this parameter value. This
parameter lists all the aggregates that are in the precommit phase of a
Copy-Free Transition workflow.

``[-is-transition-out-of-space {true|false}]`` - Is a 7-Mode transitioning
aggregate that is not yet committed in clustered Data ONTAP and is
currently out of space [.advaced-priv]#(privilege: advanced)#::
Selects the aggregates that match this parameter value. This parameter is
used to list all the 7-mode transitioning aggregates that are not yet
committed in ONTAP, and are currently out of space.

``[-autobalance-unbalanced-threshold-percent <integer>]`` - Threshold When
Aggregate Is Considered Unbalanced (%) [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the aggregates that have the specified unbalanced threshold
percentage.

``[-autobalance-available-threshold-percent <integer>]`` - Threshold When
Aggregate Is Considered Balanced (%) [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the aggregates that have the specified available threshold
percentage.

``[-resync-priority {high(fixed)|high|medium|low}]`` - Resynchronization
Priority::
This parameter indicates the relative priority that is used to decide
whether a mirrored aggregate can start a resynchronization operation or
not. This field is not set for unmirrored aggregates.
```

```
+
Use the xref:{relative_path}storage-aggregate-resynchronization-
modify.html[storage aggregate resynchronization modify] command to modify
this field for mirrored aggregates.
+
The valid values for this field are:
+
  * high(fixed): This value is reserved for ONTAP system aggregates, which
cannot have any other value for this field. It cannot be explicitly set on
a data aggregate. These aggregates always start their resynchronization
operation at the first available opportunity.
  * high: Mirrored data aggregates with this priority value start
resynchronization first.
  * medium: Mirrored data aggregates with this priority value start
resynchronization after all the system aggregates and data aggregates with
'high' priority value have started resynchronization.
  * low: Mirrored data aggregates with this priority value start
resynchronization only after all the other aggregates have started
resynchronization.



``[-data-compaction-space-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space
Saved by Data Compaction::
This parameter indicates the amount of the space saved by Data Compaction
in bytes.

``[-data-compaction-space-saved-percent <percent>]`` - Percentage Saved by
Data Compaction::
This parameter indicates the percentage of space saved in the aggregate by
Data Compaction.

``[-data-compacted-count {<integer>[KB|MB|GB|TB|PB]}]`` - Amount of
compacted data::
This parameter indicates the number of bytes occupied by compacted data
inside this aggregate.

``[-creation-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Timestamp of Aggregate
Creation::
This parameter indicates the date and time the aggregate was created.

``[-single-instance-data-logging {off|on}]`` - Enable SIDL::
If this parameter is specified, the command displays whether Single
Instance Data Logging feature is enabled on the specified aggregate.

``[-composite {true|false}]`` - Composite::
```

```
 If this parameter is specified, the command displays information only
about aggregates whose classification as a FabricPool matches the
specified value. A FabricPool has an external capacity tier attached to
it.

``[-is-fabricpool-mirrored {true|false}]`` - Is FabricPool Mirrored::
 If this parameter is specified, the command displays information only
about FabricPools whose classification as mirrored matches the specified
value. A mirrored FabricPool has a second external capacity tier attached
to it.

``[-composite-capacity-tier-used {<integer>[KB|MB|GB|TB|PB]}]`` - Capacity
Tier Used Size::
If this parameter is specified, the command displays the amount of space
in use in the attached external capacity tier.

``[-sis-space-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved by
Storage Efficiency::
This parameter indicates the total amount of space saved by storage
efficiency in bytes.

``[-sis-space-saved-percent <percent>]`` - Percentage of Space Saved by
Storage Efficiency::
This parameter indicates the percentage of space saved by storage
efficiency.

``[-sis-shared-count {<integer>[KB|MB|GB|TB|PB]}]`` - Amount of Shared
bytes count by Storage Efficiency::
This parameter indicates the number of bytes shared by storage efficiency.

``[-is-inactive-data-reporting-enabled {true|false}]`` - Inactive Data
Reporting Enabled::
If this parameter is specified, the command displays whether reporting of
inactive user data is enabled. This parameter is not allowed on
FabricPools

``[-inactive-data-reporting-start-timestamp <MM/DD/YYYY HH:MM:SS>]`` -
Timestamp when Inactive Data Reporting was Enabled::
If this parameter is specified, the command displays the timestamp at
which inactive data reporting was enabled on the aggregate. This parameter
is not allowed on FabricPools.

``[-encrypt-with-aggr-key {true|false}]`` - Enable Aggregate level
Encryption::
Selects the aggregates that are encrypted with aggregate keys.

``[-drive-protection-enabled {true|false}]`` - Aggregate uses data
```

```
protected SEDs::
If this parameter is specified, the command displays whether this
aggregate is entirely composed of self-encrypting drives that have data
protection enabled.

``[-azcs-read-optimization {on|off}]`` - azcs read optimization::
If this parameter is specified, the command displays whether azcs-with-
compression feature is enabled.

``[-space-required-for-revert {<integer>[KB|MB|GB|TB|PB]}]`` - Metadata
Reserve Space Required For Revert [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the additional amount
of free space needed for metadata reserve in the aggregate if the system
is required to revert to an ONTAP release earlier than 9.11.1x.

== Examples

The following example displays information about all aggregates that are
owned by nodes in the local cluster:
```

```
cluster1::> storage aggregate show
Aggregate Size Available Used% State #Vols Nodes RAID Status
--------- -------- --------- ----- ------- ------ ---------------- ------------
aggr0 6.21TB 1.78TB 71% online 49 cluster1-01 raid_dp,
normal
aggr1 56.04MB 55.89MB 0% online 0 cluster1-02 raid_dp,
mirrored,
normal
aggr2 1.77TB 1.63TB 8% online 1 cluster1-01 raid_dp,
normal
aggr3 1.77TB 1.73TB 2% online 2 cluster1-02 raid_dp,
normal
4 entries were displayed.
```

```
The following example displays information about an aggregate name aggr1:
```

```
cluster1::> storage aggregate show -aggregate aggr1
Aggregate: aggr1
Checksum Style: block
Number Of Disks: 6
Mirror: true
Nodes: cluster1-02
Disks for First Plex: 1.1.2,
1.1.10,
1.1.11
Disks for Mirrored Plex: 1.1.6,
1.1.8,
1.1.9
```

Free Space Reallocation: off
HA Policy: sfo
Space Reserved for Snapshot Copies: 5%
Hybrid Enabled: false
Available Size: 53.10MB
Block Type: 64-bit
Checksum Enabled: true
Checksum Status: active
Cluster: cluster1
Home Cluster ID: 686964a0-2172-11e3-837d-123478563412
DR Home ID: -
DR Home Name: -
Has Mroot Volume: false
Has Partner Node Mroot Volume: false
Home ID: 4050409551
Home Name: cluster1-02
Total Hybrid Cache Size: 0B
Hybrid: false
Inconsistent: false
Is Aggregate Home: true
Max RAID Size: 16
Hybrid Aggregate SSD Tier Maximum RAID Group Size: -
Owner ID: 4050409551
Owner Name: cluster1-02
Used Percentage: 0%
Plexes: /aggr1/plex0, /aggr1/plex1
RAID Groups: /aggr1/plex0/rg0 (block)
/aggr1/plex1/rg0 (block)
RAID Status: raid_dp, mirrored, normal
RAID Type: raid_dp
SyncMirror Resync Snapshot Frequency in Minutes: 60
Is Root: false
Space Used By metadata for Volume Efficiency: 0B
Size: 53.24MB
SnapLock Type of the Aggregate: -
State: online
Used Size: 144KB
Number Of Volumes: 0
Is Flash Pool Caching: -
Is Eligible for Auto Balance Aggregate: false
State of the aggregate being balanced: ineligible
State Change Counter for Auto Balancer: 0
Is Encrypted: true
Encryption Key ID: 40004FE300000000030300000000000436F5DB53445FD603FB5A8A64937AA7B
Is in the precommit phase of Copy-Free Transition: false
Is a 7-Mode transitioning aggregate that is not yet committed in clustered ONTAP and is currently out of space: false
Threshold When Aggregate Is Considered Unbalanced (%): 70
Threshold When Aggregate Is Considered Balanced (%): 40
Resynchronization Priority: -
Space Saved by Data Compaction: 99.24MB
Percentage Saved by Data Compaction: 7%
Amount of compacted data: 99.24MB
Timestamp of Aggregate Creation: 1/3/2017 23:38:06

Enable SIDL: off
Composite: false
Capacity Tier Used Size: 0B
Space Saved by Storage Efficiency: 99.24MB
Percentage of Space Saved by Storage Efficiency: 7%
Amount of Shared bytes count by Storage Efficiency: 99.24MB

```
The following example displays information about aggregates that are owned
by nodes in cluster1:
```

cluster1::> storage aggregate show -cluster cluster1

cluster1:
Aggregate Size Available Used% State #Vols Nodes RAID Status
--------- -------- --------- ----- ------- ------ ---------- ----------
aggr0 6.04GB 3.13GB 48% online 2 cluster1-01 raid_dp,
mirrored,
normal
aggr1 53.24MB 12.59MB 76% online 2 cluster1-02 raid_dp,
mirrored,
normal
2 entries were displayed.

```
The following example displays information about aggregates that are owned
by nodes in the remote cluster named cluster2:
```

cluster1::> storage aggregate show -cluster cluster2

cluster2:
Aggregate Size Available Used% State #Vols Nodes RAID Status
--------- -------- --------- ----- ------- ------ ---------- ----------
aggr2 - - - remote_cluster
- - -
aggr3 - - - remote_cluster
- - -
2 entries were displayed.

```
The following example displays information about aggregates that are owned
by nodes in all the clusters:
```

cluster1::> storage aggregate show -cluster *

cluster2:
Aggregate Size Available Used% State #Vols Nodes RAID Status
--------- -------- --------- ----- ------- ------ ---------- ----------
aggr2 - - - remote_cluster
- - -
aggr3 - - - remote_cluster
- - -

cluster1:
Aggregate Size Available Used% State #Vols Nodes RAID Status
--------- -------- --------- ----- ------- ------ ---------- ------------
aggr0 6.04GB 3.14GB 48% online 2 cluster1-01 raid_dp,
mirrored,
normal
aggr1 53.24MB 12.59MB 76% online 2 cluster1-02 raid_dp,
mirrored,
normal
4 entries were displayed.

```
== Related Links

* xref:{relative_path}storage-aggregate-show-space.html[storage aggregate
show-space]
* xref:{relative_path}storage-aggregate-resynchronization-
modify.html[storage aggregate resynchronization modify]




[[IDb5bdf1f62201d2965191e0442024a997]]
= storage aggregate verify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Verify an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate verify` command verifies the two plexes of an
aggregate. It compares the data in the two plexes to ensure that the
plexes are identical. It can be used whenever the administrator needs to
ensure that the two plexes are completely synchronized with each other. To
view any discrepancies, use the following command:
```

event log show -message-name raid.mirror.verify.mismatch

```
== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the aggregate to be verified. If no aggregate is
specified then the action specified by the parameter `-action` will be
taken on all the aggregates.

``-action {start|stop|resume|suspend|status}`` - Action::
This parameter specifies the action to be taken. The possible actions are:
+
  * start - Starts a verify.
  * stop - Permanently stops a verify. A stopped verify cannot be resumed.
  * resume - Resumes a suspended verify.
  * suspend - Suspends a verify.
  * status - Displays the current status of a verify.



``[-plex-to-fix <text>]`` - Plex to be Corrected in Case of Mismatches::
This parameter specifies the name of a plex to fix in case the two plexes
of the aggregate do not match. The default behavior is to log any
discrepancies instead of fixing them.
+
NOTE: This parameter is only applicable when the command is used to start
a verify.
+


== Examples

The following example starts a verify on an aggregate named aggr1.
```

cluster1::> storage aggregate verify -aggregate aggr1 -action start

```
The following example queries the status of a verify on an aggregate named
aggr1.
```

cluster1::> storage aggregate verify -aggregate aggr1 -action status
Aggregate:aggr1, Is Suspended:false, Percentage Completed:19.03%

```
The following example starts a verify on all the aggregates.
```

cluster1::> storage aggregate verify -action start

```
[[ID6699c2bb8400addbcc191212ac7cc804]]
= storage aggregate efficiency show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate storage efficiency details

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate efficiency show` command displays information about
the different storage efficiency of all the aggregates.If no parameters
are specified, the command displays the following information for all
aggregates:

* Aggregate
* Node
* Cross-vol-background-dedupe State (Enabled, Disabled)
* Cross-vol-inline-dedupe State (Enabled, Disabled)



== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
Displays the aggregate name. If this parameter is specified, the command
displays detailed information about the storage efficiency of the
specified aggregate.

``[-node {<nodename>|local}]`` - Node::
Displays the node which owns the aggregate. If this parameter is
specified, the command displays storage efficiency information only about
the aggregates that are located on the specified node.
```

```
``[-cross-volume-background-dedupe {true|false}]`` - Cross Volume
Background Deduplication::
Displays whether the cross volume background deduplication is
enabled/disabled in the aggregate.

``[-cross-volume-inline-dedupe {true|false}]`` - Cross Volume Inline
Deduplication::
Displays whether the cross volume inline deduplication is enabled/disabled
in the aggregate.

``[-cross-volume-dedupe-savings {true|false}]`` - Has Cross Volume
Deduplication Savings::
Displays whether the aggregate may have savings from cross volume
deduplication.

``[-auto-adaptive-compression-savings {true|false}]`` - Has Auto Adaptive
Compression Savings::
Displays whether the aggregate may have auto adaptive compression savings.

== Examples

The following example displays information about all aggregates that are
owned by nodes in the local cluster:
```

cluster:::> storage aggregate efficiency show

Aggregate: aggr0
Node: vivek6-vsim2

Has Cross Volume Deduplication Savings: false
Cross Volume Background Deduplication: false
Cross Volume Inline Deduplication: false

Aggregate: aggr1
Node: vivek6-vsim2

Has Cross Volume Deduplication Savings: true
Cross Volume Background Deduplication: true
Cross Volume Inline Deduplication: true
2 entries were displayed.

```
[[ID8b8f0287026385887c41aedb3eb01ae1]]
= storage aggregate efficiency stat
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display aggregate storage efficiency statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate efficiency stat` command displays storage
efficiency statistics.

== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
If this parameter is specified, the command displays storage efficiency
statistics for the specified aggregate.

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays storage efficiency
statistics for the aggregates that are located on the specified node.

``[-total-blocks-recompressed <integer>]`` - Total Number of Blocks
Recompressed::
This field represents the total number of blocks recompressed. If this
parameter is specified, the command displays storage efficiency statistics
that match the specified block count.

``[-active-blocks-recompressed <integer>]`` - Number of Active Blocks
Recompressed::
This field represents the total number of hot blocks recompressed. If this
parameter is specified, the command displays storage efficiency statistics
that match the specified block count.

``[-inactive-blocks-recompressed <integer>]`` - Number of Inactive Blocks
Recompressed::
This field represents the total number of cold blocks recompressed. If
this parameter is specified, the command displays storage efficiency
```

```
statistics that match the specified block count.

== Examples

The following example displays storage efficiency statistics for all the
aggregates that are owned by nodes in the local cluster:
```

cluster1:::> storage aggregate efficiency stat
Aggregate: aggr1
Node: node1
Total Number of Blocks Recompressed: 512
Number of Active Blocks Recompressed: 120
Number of Inactive Blocks Recompressed: 392
Aggregate: aggr2
Node: node1
Total Number of Blocks Recompressed: 0
Number of Active Blocks Recompressed: 0
Number of Inactive Blocks Recompressed: 0

2 entries were displayed.

```
[[ID5504a8f24685a55568c310f355490ed2]]
= storage aggregate efficiency cross-volume-dedupe revert-to
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reverts the cross volume deduplication savings on an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate cross-volume-dedupe revert-to` command is used to
revert cross volume deduplication savings on an aggregate.

== Parameters

``-aggregate <aggregate name>`` - Aggregate [.advaced-priv]#(privilege:
advanced)#::
This specifies the aggregate on which cross volume deduplication savings
should be reverted. If no aggregate is specified then it will revert the
savings on all aggregates

``[-c, -clean-up {true|false}]`` - Delete Previously Downgraded Metafiles
[.advaced-priv]#(privilege: advanced)#::
This specifies whether downgrade metafile needs to be removed so that
other efficiency operations can start on that aggregate.

== Examples

The following example displays information for reverting cross volume
background deduplication on aggregate "aggr1":
```

cluster:::> storage aggregate efficiency cross-volume-dedupe revert-to -aggregate aggr1
The revert operation started on aggregate "aggr1".

cluster:::> storage aggregate efficiency cross-volume-dedupe revert-to -aggregate aggr1 -clean-up true
The revert operation started on aggregate "aggr1".

```
[[ID975ed3a22a5e03d151654a9b553b8b7f]]
= storage aggregate efficiency cross-volume-dedupe show
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate cross volume deduplication efficiency details

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate efficiency cross-volume-dedupe show` command
displays information in detail about the different storage efficiency of
all the aggregates.If no parameters are specified, the command displays
the following information for all aggregates:

== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
Displays the aggregate name. If this parameter is specified, the command
displays detailed information about the storage efficiency of the
specified aggregate.

``[-node {<nodename>|local}]`` - Node::
Displays the node which owns the aggregate. If this parameter is
specified, the command displays storage efficiency information only about
the aggregates that are located on the specified node.

``[-background-progress <text>]`` - Progress::
Displays the information for the aggregates that match the specified
progress.

``[-background-op-status <text>]`` - Operation Status::
Displays the information for the aggregates that match the specified
operation status.

``[-background-last-op-state <text>]`` - Last Operation State::
```

```
Displays the information for the aggregates that match the specified last
operation state.

``[-background-last-success-op-begin <Date>]`` - Last Success Operation
Begin Time::
Displays the information for the aggregates that match the specified last
successful operation begin time.

``[-background-last-success-op-end <Date>]`` - Last Success Operation End
Time::
Displays the information for the aggregates that match the specified last
successful operation end time.

``[-background-last-op-begin <Date>]`` - Last Operation Begin Time::
Displays the information for the aggregates that match the specified last
operation begin time.

``[-background-last-op-end <Date>]`` - Last Operation End Time::
Displays the information for the aggregates that match the specified last
operation end time.

``[-background-last-op-error <text>]`` - Last Operation Error::
Displays the information for the aggregates that match the specified last
operation error.

``[-background-stage <text>]`` - Stage::
Displays the information for the aggregates that match the specified
stage.

``[-background-checkpoint-time <Date>]`` - Checkpoint Time::
Displays the information for the aggregates that match the specified
checkpoint time.

``[-background-checkpoint-op-type <text>]`` - Checkpoint Operation Type::
Displays the information for the aggregates that match the specified
checkpoint operation type.

``[-background-checkpoint-stage <text>]`` - Checkpoint Stage::
Displays the information for the aggregates that match the specified
checkpoint stage.

``[-background-dedupe {true|false}]`` - Background State::
Displays the information for the aggregates that match the specified cross
volume background dedupe state.

``[-inline-dedupe {true|false}]`` - Inline State::
Displays the information for the aggregates that match the specified cross
```

```
volume inline dedupe state.

``[-dedupe-savings {true|false}]`` - Has Cross Volume Deduplication
Savings::
Displays the information for the aggregates that has some savings from
cross volume deduplication.

== Examples

The following example displays information about all aggregates that are
owned by nodes in the local cluster:
```

cluster:::> storage aggregate efficiency cross-volume-dedupe show

Aggregate: aggr0
Node: vivek6-vsim2

Has Cross Volume Deduplication Savings: false

--------:Cross Volume Background Deduplication Status:--------
State: false
Progress: -
Operation Status: Idle
Last Operation State: Success
Last Success Operation Begin Time: -
Last Success Operation End Time: -
Last Operation Begin Time: -
Last Operation End Time: -
Last Operation Error: Operation succeeded
Stage: -
Checkpoint Time: -
Checkpoint Operation Type: -
Checkpoint Stage: -

----------:Cross Volume Inline Deduplication Status:---------
State: false
Aggregate: aggr1
Node: vivek6-vsim2

Has Cross Volume Deduplication Savings: true

--------:Cross Volume Background Deduplication Status:--------
State: true
Progress: -
Operation Status: Idle
Last Operation State: Success
Last Success Operation Begin Time: Wed Aug 30 06:31:50 2017
Last Success Operation End Time: Wed Aug 30 06:31:50 2017
Last Operation Begin Time: Wed Aug 30 06:31:50 2017
Last Operation End Time: Wed Aug 30 06:31:50 2017
Last Operation Error: Operation succeeded
Stage: Cross volume sharing Done

Checkpoint Time: -
Checkpoint Operation Type: -
Checkpoint Stage: -

----------:Cross Volume Inline Deduplication Status:---------
State: true

2 entries were displayed.

```
[[ID4f3c7cc44eed6498f003201e098e4f6c]]
= storage aggregate efficiency cross-volume-dedupe start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Starts the cross volume background deduplication on an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate efficiency cross-volume-dedupe start` command is
used to start cross volume background deduplication on an aggregate.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This specifies the aggregate on which cross volume background
deduplication should be started. If no aggregate is specified then it will
start on all aggregates

``[-s, -scan-old-data <true>]`` - Scan Old Data::
This option processes all the existing data on all volumes on the
aggregate. It prompts for user confirmation before proceeding. Default
value is `false` .

== Examples

The following example displays information for starting cross volume
background deduplication on aggregate "aggr1":
```

cluster:::> storage aggregate efficiency cross-volume-dedupe start -aggregate aggr1
The efficiency operation on aggregate "aggr1" has started.

cluster:::> storage aggregate efficiency cross-volume-dedupe start -aggregate aggr1 -scan-old-data true
The efficiency operation on aggregate "aggr1" has started.

```
[[ID35829e5f4b4db6559ad67fde916f446c]]
= storage aggregate efficiency cross-volume-dedupe stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stops the cross volume background deduplication on an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate cross-volume-dedupe stop` command is used to stop
cross volume background deduplication on an aggregate.

== Parameters

``-aggregate <aggregate name>`` - Aggregate [.advaced-priv]#(privilege:
advanced)#::
This specifies the aggregate on which cross volume background
deduplication should be stopped. If no aggregate is specified then it will
stop on all aggregates

== Examples

The following example displays information for stopping cross volume
background deduplication on aggregate "aggr1":
```

cluster:::> storage aggregate efficiency cross-volume-dedupe stop -aggregate aggr1
The efficiency operation on aggregate "aggr1" is being stopped.

```
[[ID191a555c9784c8be5d4552d206303c66]]
= storage aggregate encryption show-key-id
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Display encrypted aggregate information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate encryption show-key-id` command displays the key
IDs of all NAE (NetApp Aggregate Encryption) aggregates.

== 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.

``[-aggregate <text>]`` - Aggregate::
If this parameter is specified, the command displays information only
about the specific NAE (NetApp Aggregate Encryption) aggregate.

``[-aggrID <UUID>]`` - Aggregate UUID::
If this parameter is specified, the command displays the key ID of the
specified NAE (NetApp Aggregate Encryption) aggregate ID.

``[-keyid-index-zero <text>,...]`` - 0th Index Keyid::
If this parameter is specified, the command displays the 0th index key ID
of NAE (NetApp Aggregate Encryption) aggregates.




[[IDe389744915faf93e29571ab3beead8c1]]
= storage aggregate inode-upgrade resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume suspended inode upgrade

*Availability:* This command is available to _cluster_ administrators at
```

```
the _advanced_ privilege level.

== Description

The `storage aggregate inode-upgrade resume` command resumes a suspended
inode upgrade process. The inode upgrade process might have been suspended
earlier due to performance reasons.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command resumes the upgrade process of
an aggregate that is located on the specified node.

``-aggregate <aggregate name>`` - Aggregate Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the aggregate for which the inode upgrade process is to be
resumed.

== Examples

The following example resumes an aggregate upgrade process:
```

cluster1::> storage aggregate inode-upgrade resume -aggregate aggr1

```
[[ID42ad7ec03a6cbd84a2adfbb54935f5a6]]
= storage aggregate inode-upgrade show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display inode upgrade progress

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate inode-upgrade show` command displays information
about aggregates undergoing the inode upgrade process. The command output
depends on the parameter or parameters specified with the command. If no
parameters are specified, the command displays the default fields about
all aggregates undergoing the inode upgrade process. The default fields
```

```
are:

* aggregate
* status
* scan-percent
* remaining-time
* space-needed
* scanner-progress



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-aggregate <text>]`` - Aggregate Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified with the `-node` parameter, the command
displays detailed information about the specified aggregate. If only this
parameter is specified, the command displays information about all
aggregates that match the specified name.

``[-node <nodename>]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that are located on the specified node.

``[-status {pending|scanning|suspended-initalizing|suspended|cleanup-
pending|cleanup|cleanup-done|suspended-aborting|suspended-
removing|suspended-while-removing|suspended-ironing}]`` - Upgrade Status
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that match the specified inode upgrade
status.

``[-scan-percent <percent>]`` - Upgrade Scan Percent Complete [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregate or aggregates that match the specified inode upgrade
progress percentage.
```

```
``[-space-needed {<integer>[KB|MB|GB|TB|PB]}]`` - Space Needed to Complete
Upgrade [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregate or aggregates where the space needed to complete the
upgrade process matches the specified size.

``[-remaining-time <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Remaining Upgrade Time [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregate or aggregates where the remaining time to complete the
inode upgrade process matches the specified time.

``[-scanner-progress <text>]`` - Scanner Progress [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregate or aggregates where the progress of the inode upgrade
process matches the input.

== Examples

The following example displays information about all aggregates undergoing
the inode upgrade process:
```

cluster1::> storage aggregate inode-upgrade show
Aggregate Status %Complete Time Remaining Space Needed Inode Progress
--- --- --- --- --- ---
aggr0 pending 0% - 20.36MB Public : Inode 0 out of 65562
aggr1 pending 0% - 19.84MB Public : Inode 0 out of 63714

```
[[IDf11fb584592a22d41a294cf85fcc2cb6]]
= storage aggregate object-store attach
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Attach an object store to an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store attach` command attaches an object
```

```
store to an aggregate to create a FabricPool. This command requires two
parameters to create a FabricPool - an aggregate and a configuration to
attach an object-store to the aggregate. This command verifies whether the
object store is accessible through the intercluster LIF both from the node
on which the aggregate is present as well as its High Availability (HA)
partner node. The command fails if the object store is not accessible.

NOTE:  Once an object store is attached to an aggregate, it cannot be
detached.



== Parameters

``-aggregate <text>`` - Name of the Aggregate::
This parameter specifies the aggregate to which the object store must be
attached to create a FabricPool.

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter specifies the object store configuration that describes the
object store to be attached. The object store configuration has
information about object store server name, port, access credentials, and
provider type.

``[-allow-flexgroup {true|false}]`` - Allow Existing FlexGroup
Constituents in the Aggregate on Attach::
This optional parameter allows attaching object store to an aggregate
already containing FlexGroup constituents. The default value is false.
Mixing FabricPools and non-FabricPools within a FlexGroup is not
recommended. All aggregates hosting constituents of a FlexGroup should be
attached to the object store.

== Examples

The following example attaches an object store to aggregate aggr1:
```

cluster1::>storage aggregate object-store attach -aggregate aggr1 -object-store-name my-store

```
[[ID54c7956bb1a4855d7ec365c17bd3ab64]]
= storage aggregate object-store mirror
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Attaches a second object store to a FabricPool aggregate to create a
mirror

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store mirror` command attaches a second
object store to a FabricPool aggregate to create a mirror. This command
verifies whether the object store is accessible through the intercluster
LIF from both the node on which the aggregate is present and from its
high-availability (HA) partner node. The command fails if the object store
is not accessible.

== Parameters

``-aggregate <text>`` - Name of the Aggregate::
This parameter specifies the aggregate to which the object store must be
attached in order to create a mirror.

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter specifies the name of the new object store configuration to
be attached to the aggregate. The object store configuration has
information about the object store server name, port, access credentials,
and provider type.

== Examples

The following example shows how to create a mirror to aggregate aggr1:
```

cluster1::>storage aggregate object-store mirror -aggregate aggr1 -object-store-name my-store-2

```
[[ID4707f4e4d34ed70f98857a33be11df7f]]
= storage aggregate object-store modify
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify attributes of object stores attached to an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store modify` command is used to update one
or more object store parameters.

== Parameters

``-aggregate <text>`` - Aggregate Name::
This parameter identifies the aggregate to which the object store to be
modified is attached.

``-object-store-name <text>`` - ONTAP Name for this Object Store Config::
This parameter identifies the configuration name of the object store to be
modified.

``[-unreclaimed-space-threshold <percent>]`` - Threshold for Reclaiming
Unreferenced Space [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the usage threshold below which ONTAP
reclaims unused space from objects in the object store. When ONTAP writes
data to the object store, it packages multiple file system blocks into one
object. Over time, blocks stored in an object can be freed, leaving part
of the object unused. When the percentage of used blocks in an object
falls below this threshold, a background task moves the blocks which are
still used to a new object. Afterwards, ONTAP frees the original object to
reclaim the unused space. Valid values are between 0% and 99%. The default
value depends on the object store's provider type. It is 20% for
``_AWS_S3_`` , 25% for ``_Azure_Cloud_`` , 40% for ``_SGWS_`` , 20% for
``_IBM_COS_`` , 20% for ``_AliCloud_`` , 20% for ``_GoogleCloud_`` and 40%
for ``_ONTAP_S3_`` . Consult the FabricPool best practices guidelines for
more information.

``[-tiering-fullness-threshold <percent>]`` - Aggregate Fullness Threshold
Required for Tiering [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the percentage of space in the
performance tier which must be used before data is tiered out to the
capacity tier.

``[-mirror-type {primary|mirror}]`` - Object Store Mirror Type::
```

```
This parameter specifies the object store mirror type. Valid mirror types
are ``_primary_`` or ``_mirror_`` .

``[-force-tiering-on-metrocluster {true|false}]`` - Force Tiering with no
Mirror in a MetroCluster Configuration::
This parameter specifies force tiering option enabled or not on primary
object stores for aggregates in a MetroCluster configuration.

``[-migrate-threshold <percent>]`` - Minimum Aggregate Space Threshold
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the minimum percentage of performance
tier free space that must exist in order for migration of data from the
capacity tier to performance tier to be allowed.

== Examples

The following example modifies the unreclaimed space threshold of an
object store attached to an aggregate named aggr1:
```

cluster1::> storage aggregate object-store modify -aggregate aggr1 -object-store-name my-store -unreclaimed -space-threshold 20%

```
[[ID3ade6bf61fea1abe1a763ecfe6e3d032]]
= storage aggregate object-store show-freeing-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show status of background object freeing work after aggregate delete

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate object-store show-freeing-status` command displays
status information about the background work that frees an aggregate's
objects from an object store after a xref:{relative_path}storage-
aggregate-delete.html[storage aggregate delete] .

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-bin-uuid <UUID>]`` - UUID of the Bin [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the aggregate attached to the specified bin UUID.

``[-config-id <integer>]`` - Object Store Config ID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the aggregate attached to the object- store with specified config
ID.

``[-object-store-name <text>]`` - Object Store Configuration Name
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about object stores whose configuration name matches the specified names.

``[-aggregate-name <aggregate name>]`` - Aggregate [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the specified aggregates that were deleted.

``[-request-state {queued|running|cleaning-up|finishing}]`` - Request
State [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the object stores that have the specified object freeing request
state.

``[-num-objects-freed <integer>]`` - Num Objects Freed [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the object stores that have the specified number of objects that
have been freed.

``[-last-error <text>]`` - The Last Error Encountered [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the object stores that have the specified last error encountered.

== Related Links
```

```
* xref:{relative_path}storage-aggregate-delete.html[storage aggregate
delete]




[[ID69128f51dc299ff3b3c4022a5010ddf8]]
= storage aggregate object-store show-resync-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display object store mirror resync progress

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store show-resync-status` command displays
information about the resync progress for mirror object stores attached to
a FabricPool.

== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
If this parameter is specified, the command displays information only
about the specified aggregates.

``[-primary-object-store-name <text>]`` - Primary Object Store Name::
If this parameter is specified, the command displays information only
about aggregates whose primary object store configuration name matches the
specified names.

``[-mirror-object-store-name <text>]`` - Mirror Object Store Name::
If this parameter is specified, the command displays information only
```

```
about aggregates whose mirror object store configuration name matches the
specified names.

``[-percent-complete <percent>]`` - Resync Complete Percentage::
If this parameter is specified, the command displays information only
about aggregates whose mirror object store resync progress percentage
matches the specified value.

== Examples

The following example displays resync progress for all aggregates with
mirror object stores which are not in sync:
```

cluster1::>storage aggregate object-store show-resync-status

```
[[ID46242c695db05d362b2c245b99e22015]]
= storage aggregate object-store show-space
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display space utilization of object stores attached to an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The ` storage aggregate object-store show-space` command displays
information about the amount of space used in the object store for each of
the aggregates in FabricPool. The used space is displayed in both absolute
size as well as a percentage of the FabricPool license limit.

== 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.
```

```
``[-aggregate <text>]`` - Aggregate Name::
If this parameter is specified, the command displays space information
only about object stores that are attached to the specified aggregates.

``[-object-store-name <text>]`` - ONTAP Name for this Object Store
Config::
If this parameter is specified, the command displays space information
only about object stores whose configuration name matches the specified
names.

``[-object-store-availability <object Store Availability>]`` -
Availability of the Object Store::
If this parameter is specified, the command displays space information
about the object stores whose availability status matches the specified
value. Supported values with this parameter are - available and
unavailable.

``[-object-store-unavailable-reason <text>]`` - Reason why Object Store is
Unavailable::
If this parameter is specified, the command displays information only
about object stores whose unavailability reason matches the specified
reason.

``[-provider-type <providerType>]`` - Type of the Object Store Provider::
If this parameter is specified, the command displays information only
about object store configurations whose provider type matches the
specified value.

``[-license-used-percent <percent_no_limit>]`` - License Space Used
Percent::
If this parameter is specified, the command displays space information
only about object stores whose space used by the associated aggregate as a
percentage of the license limit matches the specified value. If the object
store does not require a license, then this field is not set.

``[-unreclaimed-space-threshold <percent>]`` - Threshold for Reclaiming
Unreferenced Space [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about object stores whose threshold for reclaiming unused space from
objects in the object store matches the specified value.

``[-tiering-fullness-threshold <percent>]`` - Aggregate Fullness Threshold
Required for Tiering [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about object stores whose performance tier fullness threshold for tiering
matches the specified value.
```

```
``[-mirror-type {primary|mirror}]`` - Object Store Mirror Type::
If this parameter is specified, the command displays information about
object stores with a mirror-type that matches the specified value.

``[-is-mirror-degraded {true|false}]`` - This object store is in mirror
degraded mode::
If this parameter is specified, the command displays information only
about mirror object stores which have the specified mirror degraded state.

``[-force-tiering-on-metrocluster {true|false}]`` - Force Tiering with no
Mirror in a MetroCluster Configuration::
If this parameter is specified, the command displays information only
about primary object stores for which force tiering is toggled on for
aggregates in a MetroCluster configuration.

``[-cluster <Cluster name>]`` - The name of the Cluster to which the bin
belongs::
If this parameter is specified, the command displays information only
about object stores for which cluster matches the specified value.

``[-migrate-threshold <percent>]`` - Minimum Aggregate Space Threshold
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about object stores whose performance tier migrate threshold matches the
specified value.

== Examples

The following example displays space information about all object stores:
```

cluster1::>storage aggregate object-store show-space

```
[[ID2524fb6e3a0f5cbb896bb7474ed35a46]]
= storage aggregate object-store show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the details of object stores attached to an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `storage aggregate object-store show` command displays information
about all the object stores in the system.

== 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.

``[-aggregate <text>]`` - Aggregate Name::
If this parameter is specified, the command displays information only
about the object stores that are attached to the specified aggregates.

``[-object-store-name <text>]`` - ONTAP Name for this Object Store
Config::
If this parameter is specified, the command displays information only
about object stores whose configuration name matches the specified names.

``[-object-store-availability <object Store Availability>]`` -
Availability of the Object Store::
If this parameter is specified, the command displays information only
about object stores whose availability status matches the specified value.
Supported values with this parameter are - ``_available_`` and
``_unavailable_`` .

``[-object-store-unavailable-reason <text>]`` - Reason why Object Store is
Unavailable::
If this parameter is specified, the command displays information only
about object stores whose unavailability reason matches the specified
reason.

``[-provider-type <providerType>]`` - Type of the Object Store Provider::
If this parameter is specified, the command displays information only
about object store configurations whose provider type matches the
specified value.

``[-license-used-percent <percent_no_limit>]`` - License Space Used
Percent::
If this parameter is specified, the command displays information only
about object stores whose space used by the aggregate as a percentage of
```

```
the license limit matches the specified value.

``[-unreclaimed-space-threshold <percent>]`` - Threshold for Reclaiming
Unreferenced Space [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about object stores whose threshold for reclaiming unused space from
objects in the object store matches the specified value.

``[-tiering-fullness-threshold <percent>]`` - Aggregate Fullness Threshold
Required for Tiering [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about object stores whose performance tier fullness threshold for tiering
matches the specified value.

``[-mirror-type {primary|mirror}]`` - Object Store Mirror Type::
If this parameter is specified, the command displays information about
object stores with a mirror-type that matches the specified value.

``[-is-mirror-degraded {true|false}]`` - This object store is in mirror
degraded mode::
If this parameter is specified, the command displays information only
about mirror object stores which have the specified mirror degraded state.
When a mirror object store is attached to a FabricPool, it is initially
degraded because it does not contain a copy of all the data in the primary
object store. While the mirror is degraded, all reads are served from the
primary object store, and the mirror cannot be promoted to become the
primary. After the resync process copies all data from the primary object
store to the mirror, the mirror is no longer degraded. From that point on
the mirror is always kept in sync with the primary and never becomes
degraded again.

``[-force-tiering-on-metrocluster {true|false}]`` - Force Tiering with no
Mirror in a MetroCluster Configuration::
If this parameter is specified, the command displays information only
about primary object stores for which force tiering is toggled on for
aggregates in a MetroCluster configuration.

``[-cluster <Cluster name>]`` - The name of the Cluster to which the bin
belongs::
If this parameter is specified, the command displays information only
about object stores for which cluster matches the specified value.

``[-migrate-threshold <percent>]`` - Minimum Aggregate Space Threshold
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about object stores whose performance tier migrate threshold matches the
```

```
specified value.

== Examples

The following example displays all information about all object stores:
```

cluster1::>storage aggregate object-store show

```
[[IDcc9aa96ad8f5b998ed253ac377c3bd1e]]
= storage aggregate object-store unmirror
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove the second object store from a mirrored FabricPool

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store unmirror` command detaches a second
object store from a mirrored FabricPool. This command verifies whether the
primary object store is accessible through the intercluster LIF from both
the node on which the aggregate is present and from its high-availability
(HA) partner node. The command fails if the primary object store is not
available.

== Parameters

``-aggregate <text>`` - Name of the Aggregate::
This parameter specifies the aggregate from which the mirrored object
store must be removed.

== Examples

The following example shows how to unmirror a mirrored FabricPool aggr1:
```

cluster1::>storage aggregate object-store unmirror -aggregate aggr1

```
[[IDfc1f94f58104a7c2e44f570cd33ec60e]]
= storage aggregate object-store config create
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Define the configuration for an object store

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store config create` command is used by a
cluster administrator to tell ONTAP how to connect to an object store.
Following pre-requisites must be met before creating an object store
configuration in ONTAP.

*  A valid data bucket or container must be created with the object store
provider. This assumes that the user has valid account credentials with
the object store provider to access the data bucket.
*  The ONTAP node must be able to connect to the object store. This
includes

  *  Fast, reliable connectivity to the object store.
  *  An inter-cluster LIF (Logical Interface) must be configured on the
cluster. ONTAP will verify connectivity prior to saving this configuration
information.
  *  If SSL/TLS authentication is required, then valid certificates must
be installed.
  *  FabricPool license (required for Amazon S3 and Azure object stores).



An object-store configuration once created must not be reassociated with a
different object-store or container. See xref:{relative_path}storage-
aggregate-object-store-config-modify.html[storage aggregate object-store
config modify] command for more information. If neither the `access-key`
nor the `secret-password` are provided while setting up a configuration
for AWS_S3 object store in Cloud Volumes ONTAP, then the access key
(access key ID), the secret password (secret access key), and the session
token will be retrieved from EC2 instance metadata for the AWS Identity
and Access Management (IAM) role associated with the EC2 instance. If
ONTAP is unable to create a object store configuration, then the command
will fail explaining the reason for failure.

== Parameters
```

```
``-object-store-name <text>`` - Object Store Configuration Name::
This parameter specifies the name that will be used to identify the object
store configuration. The name can contain the following characters: "_",
"-", A-Z, a-z, and 0-9. The first character must be one of the following:
"_", A-Z, or a-z.

``-provider-type <providerType>`` - Type of the Object Store Provider::
This parameter specifies the type of object store provider that will be
attached to the aggregate. Valid options are: AWS_S3 (Amazon S3 storage),
Azure_Cloud (Microsoft Azure Cloud), SGWS (StorageGrid WebScale), IBM_COS
(IBM Cloud Object Storage), AliCloud (Alibaba Cloud Object Storage
Service), GoogleCloud (Google Cloud Storage) and ONTAP_S3.

``[-auth-type <object_store_auth_type>]`` - Authentication Used to Access
the Object Store::
This parameter specifies where the system obtains credentials for
authentication to an object store. The available choices depend on the
platform (Cloud Volumes ONTAP or not) and provider-type (AWS_S3 or not).
The ``_keys_`` value is always applicable, and if selected means that the
`access-key` and `secret-password` are provided by the system
administrator. In Cloud Volumes ONTAP, the ``_EC2-IAM_`` value is also
applicable. It means that the IAM role is associated with the EC2
instance, and that the `access-key` , `secret-password` and session token
are retrieved from EC2 instance metadata for this IAM role. Note that `-
use-iam-role` and `-auth-type` are mutually exclusive, `-auth-type EC2-
IAM` is an equivalent of `-use-iam-role true` , and `-auth-type key` is an
equivalent of `-use-iam-role false` . In Cloud Volumes ONTAP, the ``_GCP-
SA_`` value may also be applicable. It means that a session token is
retrieved from the GCP instance metadata for the Service Account
associated with the GCP instance. Similarly, ``_Azure-MSI_`` means that a
session token is retrieved from the Azure instance metadata for the
Managed Service Identity (MSI) associated with the Azure instance. For the
``_AWS_S3_`` provider, the ``_CAP_`` (C2S Authentication Portal) value is
also applicable. This should only be used when accessing C2S (Commercial
Cloud Services). If the ``_CAP_`` value is specified, then the`-cap-url`
must be specified. See `cap-url` .

``[-cap-url <text>]`` - URL to Request Temporary Credentials for C2S
Account::
This parameter is available only when `-auth-type` is ``_CAP_`` . It
specifies a full URL of the request to a CAP server for retrieving
temporary credentials (access-key, secret-pasword and session token) for
accessing the object store server. The CAP URL may look like:
`https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency=mymissio
n=myrole`
```

```
``-server <Remote InetAddress>`` - Fully Qualified Domain Name of the
Object Store Server::
This parameter specifies the Fully Qualified Domain Name (FQDN) of the
remote object store server. For Amazon S3, server name must be an AWS
regional endpoint in the format s3.amazonaws.com or s3-
<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The
region of the server and the bucket must match. For more information on
AWS regions, refer to 'Amazon documentation on AWS regions and endpoints'.
For Azure, if the `-server` is a "blob.core.windows.net" or a
"blob.core.usgovcloudapi.net", then a value of `-azure-account` followed
by a period will be added in front of the server.

``[-is-ssl-enabled {true|false}]`` - Is SSL/TLS Enabled::
This parameter indicates whether a secured SSL/TLS connection will be used
during data access to the object store. The default value is `true` .

``[-port <integer>]`` - Port Number of the Object Store::
This parameter specifies the port number on the remote server that ONTAP
will use while establishing connection to the object store.

``-container-name <text>`` - Data Bucket/Container Name::
This parameter specifies the data bucket or container that will be used
for read and write operations.
+
NOTE: This name cannot be modified once a configuration is created.
+


{ ``[-access-key <text>]`` - Access Key ID for S3 Compatible Provider
Types::
This parameter specifies the access key (access key ID) required to
authorize requests to the AWS S3, SGWS, IBM COS object stores and
ONTAP_S3. For an Azure object store see `-azure-account` .

``[-secret-password <text>]`` - Secret Access Key for S3 Compatible
Provider Types::
This parameter specifies the password (secret access key) to authenticate
requests to the AWS S3, SGWS, IBM COS object stores and ONTAP_S3. If the
`-access-key` is specified but the `-secret-password` is not, then one
will be asked to enter the `-secret-password` without echoing the input.
For an Azure object store see `-azure-private-key` .

| ``[-azure-account <text>]`` - Azure Account::
This parameter specifies the account required to authorize requests to the
Azure object store. For other object store providers see access-key.
+
```

```
NOTE: The value of this field cannot be modified once a configuration is
created.
+


``[-ask-azure-private-key {true|false}]`` - Ask to Enter the Azure Access
Key without Echoing::
If this parameter is true then one will be asked to enter `-azure-private-
key` without echoing the input. Default value: ``_true_`` .

``[-azure-private-key <text>]`` - Azure Access Key }::
This parameter specifies the access key required to authenticate requests
to the Azure object store. See also `ask-azure-private-key` . For other
object store providers see `-secret-password` .

``[-azure-sas-token <text>]`` - Azure Account Shared Access Signature
token [.advaced-priv]#(privilege: advanced)# }::
This parameter specifies the shared access signature token to authenticate
requests and provide limited access to storage resources in the Azure
object store.

``[-ipspace <IPspace>]`` - IPspace to Use in Order to Reach the Object
Store::
This optional parameter specifies the IPspace to use to connect to the
object store. Default value: ``_Default_`` .

``[-is-certificate-validation-enabled {true|false}]`` - Is SSL/TLS
Certificate Validation Enabled::
This parameter indicates whether an SSL/TLS certificate of an object store
server is validated whenever an SSL/TLS connection to an object store
server is established. This parameter is only applicable when `is-ssl-
enabled` is `true` . The default value is `true` . It is recommended to
use the default value to make sure that ONTAP connects to a trusted object
store server, otherwise identities of an object store server are not
verified.

``[-use-http-proxy {true|false}]`` - Use HTTP Proxy::
This optional parameter indicates whether an HTTP proxy will be used for
connecting to an object store. Note that an HTTP proxy is configured using
the `vserver http-proxy` commands at the `diagnostic` privilege level.
Default value: ``_false_`` .

``[-cluster <Cluster name>]`` - The Name of the Cluster to which the
Configuration Belongs::
This optional parameter should only be specified in MetroCluster switched-
over mode and specifies the name of the cluster for which the
```

```
configuration must be created. By default the configuration is created for
the local cluster.

``[-server-side-encryption {none | SSE-S3}]`` - Encryption of Data at Rest
by the Object Store Server::
This parameter specifies if AWS or other S3 compatible object store server
must encrypt data at rest. The available choices depend on provider-type.
`none` encryption (no encryption required) is supported by all types of S3
(non-Azure) object store servers. `SSE-S3` encryption is supported by and
is a default for all types of S3 (non-Azure) object store servers except
ONTAP_S3. `SSE-KMS` data at rest are encrypted by the object store server
with data keys managed by the object store server using AWS managed CMK or
customer managed CMK. `DSSE-KMS` data at rest are dual encrypted by the
object store server with data keys managed by the object store server
using AWS managed CMK or customer managed CMK. In most cases it is best
not to change default value of "SSE-S3" for object store servers which
support SSE-S3 encryption. The encryption is in addition to any encryption
done by ONTAP at a volume or at an aggregate level.

``[-url-style {path-style | virtual-hosted-style}]`` - URL Style Used to
Access S3 Bucket::
This parameter specifies the URL style used to access S3 bucket. This
option is only available for non-Azure object store providers. The
available choices and default value depend on provider-type.

``[-azure-msi-token <text>]`` - Azure MSI Token for Authentication::
This optional parameter specifies the Managed Service Identity (MSI) token
required to authenticate to the Azure object store. This form of
authentication is only supported on Azure NetApp Files.

== Examples

The following example creates an object store configuration in ONTAP:
```

cluster1::>storage aggregate object-store config create -object-store-name my_aws_store -provider-type AWS_S3 -server s3.amazonaws.com -container-name my-aws-bucket -access-key DXJRXHPXHYXA9X31X3JX

```
== Related Links

* xref:{relative_path}storage-aggregate-object-store-config-
modify.html[storage aggregate object-store config modify]
```

```
[[ID1c23985382880ccd9f0c365d1e555014]]
= storage aggregate object-store config delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete the configuration of an object store

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store config delete` command removes an
existing object store configuration in ONTAP. The configuration cannot be
deleted if it is used by any aggregates or if the system is still freeing
objects from the object store from a previously executed
xref:{relative_path}storage-aggregate-delete.html[storage aggregate
delete] command. The command xref:{relative_path}storage-aggregate-object-
store-show.html[storage aggregate object-store show] can be used to view
the aggregates attached to the object store before issuing the delete
command.

NOTE: The xref:{relative_path}storage-aggregate-object-store-
show.html[storage aggregate object-store show] command will not display
aggregates that have been previously deleted but still has objects in the
object store.



== Parameters

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter specifies the object store configuration to be deleted.

== Examples

The following example deletes an object store configuration named my-
store:
```

cluster1::>storage aggregate object-store config delete -object-store-name my-store

```
== Related Links
```

```
* xref:{relative_path}storage-aggregate-delete.html[storage aggregate
delete]
* xref:{relative_path}storage-aggregate-object-store-show.html[storage
aggregate object-store show]




[[ID45bad3c3555da8c0da4faede833e79c6]]
= storage aggregate object-store config modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify object store configuration attributes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate object-store config modify` command is used to
update one or more of object store configuration parameters. This command
must not be used to reassociate an existing valid object-store
configuration to a new object-store or container. The modifiable
parameters include object store name, server name, port, access_key,
secret-password, ipspace and is-ssl-enabled.

== Parameters

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter identifies the configuration to be modified.

``[-new-object-store-name <text>]`` - Object Store Configuration New
Name::
This optional parameter specifies the new name for the object store
configuration.

``[-auth-type <object_store_auth_type>]`` - Authentication Used to Access
the Object Store::
This optional parameter specifies where the system obtains credentials for
authentication to an object store. The available choices depend on the
platform (Cloud Volumes ONTAP or not) and provider-type (AWS_S3 or not).
The ``_keys_`` value is always applicable, and if selected means that the
`access-key` and `secret-password` are provided by the system
```

```
administrator. In Cloud Volumes ONTAP, the ``_EC2-IAM_`` value is also
applicable. It means that the IAM role is associated with the EC2
instance, and that the `access-key` , `secret-password` and session token
are are retrieved from EC2 instance metadata for this IAM role. Note that
`-use-iam-role` and `-auth-type` are mutually exclusive, `-auth-type EC2-
IAM` is an equivalent of `-use-iam-role true` , and `-auth-type key` is an
equivalent of `-use-iam-role false` . In Cloud Volumes ONTAP, the ``_GCP-
SA_`` value may also be applicable. It means that a session token is
retrieved from the GCP instance metadata for the Service Account
associated with the GCP instance. Similarly, ``_Azure-MSI_`` means that a
session token is retrieved from the Azure instance metadata for the
Managed Service Identity (MSI) associated with the Azure instance. For the
``_AWS_S3_`` provider, the ``_CAP_`` (C2S Authentication Portal) value is
also applicable. This should only be used when accessing C2S (Commercial
Cloud Services). If the ``_CAP_`` value is specified, then the`-cap-url`
must be specified. See `cap-url` .

``[-cap-url <text>]`` - URL to Request Temporary Credentials for C2S
Account::
This parameter is available only when `-auth-type` is ``_CAP_`` . It
specifies a full URL of the request to a CAP server for retrieving
temporary credentials (access-key, secret-pasword and session token) for
accessing the object store server. The CAP URL may look like:
`https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency=mymissio
n=myrole`

``[-server <Remote InetAddress>]`` - Fully Qualified Domain Name of the
Object Store Server::
This optional parameter specifies the new Fully Qualified Domain Name
(FQDN) of the same object store server. For Amazon S3, server name must be
an AWS regional endpoint in the format s3.amazonaws.com or s3-
<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The
region of the server and the bucket must match. For more information on
AWS regions, refer to 'Amazon documentation on AWS regions and endpoints'.
For Azure, if the `-server` is a "blob.core.windows.net" or a
"blob.core.usgovcloudapi.net", then the value of azure-account in the
configuration followed by a period will be added in front of the server.
Note that the value of `azure-account` cannot be modified.

``[-is-ssl-enabled {true|false}]`` - Is SSL/TLS Enabled::
This optional parameter indicates whether a secured SSL/TLS connection
will be used during data access to the object store.

``[-port <integer>]`` - Port Number of the Object Store::
This optional parameter specifies a new port number to connect to the
object store server indicated in the `-server` parameter.
```

```
``[-access-key <text>]`` - Access Key ID for S3 Compatible Provider
Types::
This optional parameter specifies a new access key (access key ID) for the
AWS S3, SGWS, IBM COS object stores and ONTAP S3.

``[-secret-password <text>]`` - Secret Access Key for S3 Compatible
Provider Types::
This optional parameter specifies a new password (secret access key) for
the AWS S3, SGWS, IBM COS object stores and ONTAP S3. For an Azure object
store see `-azure-private-key` . If the `-access-key` is specified but the
`-secret-password` is not then one will be asked to enter the `-secret-
password` without echoing the input.

``[-ask-azure-private-key {true|false}]`` - Ask to Enter the Azure Access
Key without Echoing::
If this optional parameter is true then one will be asked to enter the `-
azure-private-key` without echoing the input.

``[-azure-private-key <text>]`` - Azure Access Key::
This optional parameter specifies a new access key for Azure object store.
For other object store providers see `secret-password` . See also `ask-
azure-private-key` .

``[-azure-sas-token <text>]`` - Azure Account Shared Access Signature
token [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the shared access signature token to authenticate
requests and provide limited access to storage resources in the Azure
object store. Any previously stored token will be overwritten by the token
specified in this parameter. Pass an empty string in single quotes to
clear any previously stored token.

``[-ipspace <IPspace>]`` - IPspace to Use in Order to Reach the Object
Store::
This optional parameter specifies new ipspace values for the
configuration.

``[-is-certificate-validation-enabled {true|false}]`` - Is SSL/TLS
Certificate Validation Enabled::
This optional parameter indicates whether an SSL/TLS certificate of an
object store server is validated whenever an SSL/TLS connection to an
object store server is established. This parameter is only applicable when
`is-ssl-enabled` is `true` . It is recommended to keep the default value
which is `true` to make sure that ONTAP connects to a trusted object store
server, otherwise identities of an object store server are not verified.

``[-use-http-proxy {true|false}]`` - Use HTTP Proxy::
```

```
This optional parameter indicates whether an HTTP proxy will be used for
connecting to an object store. Note that an HTTP proxy is configured using
the `vserver http-proxy` commands at the `diagnostic` privilege level.

``[-server-side-encryption {none | SSE-S3}]`` - Encryption of Data at Rest
by the Object Store Server::
This parameter specifies if AWS or other S3 compatible object store server
must encrypt data at rest. The available choices depend on provider-type.
`none` encryption (no encryption required) is supported by all S3 (non-
Azure) object store servers. `SSE-S3` encryption is supported by all S3
(non-Azure) object store servers except ONTAP_S3. In most cases it is best
not to change default value of "SSE-S3" for object store servers which
support SSE-S3 encryption. The encryption is in addition to any encryption
done by ONTAP at a volume or at an aggregate level. Note that changing
this option does not change encryption of data which already exist in the
object store.

``[-url-style {path-style | virtual-hosted-style}]`` - URL Style Used to
Access S3 Bucket::
This parameter specifies the URL style used to access S3 bucket. This
option is only available for non-Azure object store providers. The
available choices and default value depend on provider-type.

``[-azure-msi-token <text>]`` - Azure MSI Token for Authentication::
This optional parameter specifies the new Managed Service Identity (MSI)
token required to authenticate to the Azure object store. This form of
authentication is only supported on Azure NetApp Files.

== Examples

The following example modifies two parameters (port number and is-ssl-
enabled) of an object store configuration named my-store:
```

cluster1::>storage aggregate object-store config modify -object-store-name my-store -port 1235 -is-ssl-enabled true

```
[[ID326df53340983a2a0d5b38547b202875]]
= storage aggregate object-store config rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename an existing object store configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

 The `storage aggregate object-store config rename` command is used to
rename an exiting object store configuration.

== Parameters

``-object-store-name <text>`` - Object Store Configuration Name::
This parameter identifies an existing object store configuration.

``-new-object-store-name <text>`` - Object Store Configuration New Name::
This parameter specifies the new object store configuration name.

== Examples

The following example renames an object store configuration from my-store
to ms1:
```

cluster1::>storage aggregate object-store config rename -object-store-name my-store -new-object-store-name ms1

```
[[ID1f2ec0110b303428296e1efd5aff6367]]
= storage aggregate object-store config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of object store configurations

*Availability:* This command is available to _cluster_ administrators at
```

```
the _admin_ privilege level.

== Description

The `storage aggregate object-store config show` command displays
information about all existing object store configurations in the cluster.

== 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.

``[-object-store-name <text>]`` - Object Store Configuration Name::
If this parameter is specified, the command displays information only
about object store configurations whose name matches the specified names.

``[-object-store-uuid <UUID>]`` - UUID of the Object Store Configuration::
If this parameter is specified, the command displays information only
about object store configurations whose UUID matches the specified UUID
values.

``[-provider-type <providerType>]`` - Type of the Object Store Provider::
If this parameter is specified, the command displays information only
about object store configurations whose provider type matches the
specified value.

``[-auth-type <object_store_auth_type>]`` - Authentication Used to Access
the Object Store::
If this parameter is specified, the command displays information only
about object store configurations whose authentication type matches the
specified value.

``[-cap-url <text>]`` - URL to Request Temporary Credentials for C2S
Account::
If this parameter is specified, the command displays information only
about object store configurations whose CAP URL matches the specified
value.

``[-server <Remote InetAddress>]`` - Fully Qualified Domain Name of the
Object Store Server::
If this parameter is specified, the command displays information only
```

```
about object store configurations whose server name matches the specified
value. The server name is specified as a Fully Qualified Domain Name
(FQDN).

``[-is-ssl-enabled {true|false}]`` - Is SSL/TLS Enabled::
If this parameter is specified, the command displays information only
about object store configurations whose status about the use of secured
communication over the network matches the specified value.

``[-port <integer>]`` - Port Number of the Object Store::
If this parameter is specified, the command displays information only
about object store configurations whose port numbers matches the specified
value.

``[-container-name <text>]`` - Data Bucket/Container Name::
If this parameter is specified, the command displays information only
about object store configurations whose container name matches the
specified value. ONTAP uses this container name or object store data
bucket while accessing data from the object store.

``[-access-key <text>]`` - Access Key ID for S3 Compatible Provider
Types::
If this parameter is specified, the command displays information only
about AWS S3, SGWS, IBM COS object store configurations and ONTAP S3 whose
access key matches the specified value. ONTAP requires the access key for
authorized access to the object store.

``[-azure-account <text>]`` - Azure Account::
If this parameter is specified, the command displays information only
about Azure object store configurations whose account matches the
specified value. ONTAP requires the Azure account for authorized access to
the Azure object store.

``[-ipspace <IPspace>]`` - IPspace to Use in Order to Reach the Object
Store::
If this parameter is specified, the command displays information only
about object store configurations whose IPspace matches the specified
value. ONTAP uses the IPspace value to connect to the object store.

``[-iam-role <text>]`` - IAM Role for AWS Cloud Volumes ONTAP::
If this parameter is specified, the command displays information only
about object store configurations whose IAM (Identity and Access
Management) role matches the specified value.

``[-is-certificate-validation-enabled {true|false}]`` - Is SSL/TLS
Certificate Validation Enabled::
```

```
If this parameter is specified, the command displays information only
about object store configurations whose status about the validation of
SSL/TLS certificate matches the specified value.

``[-use-http-proxy {true|false}]`` - Use HTTP Proxy::
If this parameter is specified, the command displays information only
about object store configurations for which usage of HTTP proxy matches
the specified value.

``[-cluster <Cluster name>]`` - The Name of the Cluster to which the
Configuration Belongs::
If this parameter is specified, the command displays information only
about object store configurations for which cluster matches the specified
value.

``[-server-side-encryption {none | SSE-S3}]`` - Encryption of Data at Rest
by the Object Store Server::
If this parameter is specified, the command displays information only
about object store configurations whose server-side encryption matches the
specified value.

``[-url-style {path-style | virtual-hosted-style}]`` - URL Style Used to
Access S3 Bucket::
If this parameter is specified, the command displays information only
about object store configurations whose URL style matches the specified
value.

== Examples

The following example displays all available object store configuration in
the cluster:
```

cluster1::>storage aggregate object-store config show

```
[[IDa5160bedc55dc458faf5aa562c5c0878]]
= storage aggregate object-store profiler abort
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Abort object store profiler

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate object-store profiler abort` command will abort an
ongoing object store profiler run. This command requires two parameters -
an object store configuration and a node on which the profiler is
currently running.

== Parameters

``-node {<nodename>|local}`` - Node on Which the Profiler Should Run
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the node on which the object store profiler is
running.

``-object-store-name <text>`` - Object Store Configuration Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the object store configuration that describes the
object store. The object store configuration has information about the
object store server name, port, access credentials, and provider type.

== Examples

The following example aborts the object store profiler :
```

cluster1::>storage aggregate object-store profiler abort -object-store-name my-store -node my-node

```
[[ID4ec9654a61400b75f82abe67d339cf51]]
= storage aggregate object-store profiler show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Show object store profiler status

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate object-store profiler show` command is used to
monitor progress and results of the xref:{relative_path}storage-aggregate-
object-store-profiler-start.html[storage aggregate object-store profiler
start] command.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified node.

``[-object-store-name <text>]`` - ONTAP Name for this Object Store
Configuration [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
about the objects which are placed on the specified object store.

``[-object-name-prefix <UUID>]`` - Bin UUID [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information about all
the objects whose object names are prefixed with the specified object name
prefix. Objects from two different object stores will not have the same
object name prefix.

``[-object-prefix <text>]`` - Prefix Added to Each Object [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
about the objects which have the specified prefix. Objects belonging to
two different object stores can have the same object prefix.
```

```
``[-profiler-status <text>]`` - Profiler Status [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified node.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Profiler Start Time [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified start time of the profiler run.

``[-op-name <text>]`` - Operation Name - PUT/GET [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified operation name. Possible values are PUT or
GET.

``[-op-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of Operation [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified size of the operation.

``[-op-count <integer>]`` - Number of Operations Performed [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified node.

``[-op-failed <integer>]`` - Number of Operations Failed [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified node.

``[-op-latency-minimum <integer>]`` - Minimum Latency for Operation in
Milliseconds [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified minimum latency for the operation. The
values are in milliseconds.

``[-op-latency-maximum <integer>]`` - Maximum Latency for Operation in
Milliseconds [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified maximum latency for the operation. The
values are in milliseconds.

``[-op-latency-average <integer>]`` - Average Latency for Operation in
Milliseconds [.advaced-priv]#(privilege: advanced)#::
```

```
If this parameter is specified, the command only displays information
corresponding to the specified average latency for the operation. The
values are in milliseconds.

``[-op-throughput {<integer>[KB|MB|GB|TB|PB]}]`` - Throughput per Second
for the operation [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified throughput for the operation.

``[-op-errors <text>,...]`` - Error Reasons and Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the count and specified reason for error.

``[-op-latency-histogram <text>,...]`` - Latency Histogram [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command only displays information
corresponding to the specified latency histogram for the operation.

== Examples

The following example displays the results of xref:{relative_path}storage-
aggregate-object-store-profiler-start.html[storage aggregate object-store
profiler start] :
```

cluster1::>storage aggregate object-store profiler show

```
== Related Links

* xref:{relative_path}storage-aggregate-object-store-profiler-
start.html[storage aggregate object-store profiler start]




[[ID23bd14d93a55216f0419e316072ef233]]
= storage aggregate object-store profiler start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start the object store profiler to measure latency and throughput

*Availability:* This command is available to _cluster_ administrators at
```

```
the _advanced_ privilege level.

== Description

The `storage aggregate object-store profiler start` command writes objects
to an object store and reads those objects to measure latency and
throughput of an object store. This command requires two parameters - an
object store configuration and node from which to send the PUT/GET/DELETE
operations. This command verifies whether the object store is accessible
through the intercluster LIF of the node on which it runs. The command
fails if the object store is not accessible. The command will create a
10GB dataset by doing 2500 PUTs for a maximum time period of 60 seconds.
Then it will issue GET operations of different sizes - 4KB, 8KB, 32KB,
256KB for a maximum time period of 180 seconds. Finally it will delete the
objects it created. This command can result in additional charges to your
object store account. This is a CPU intensive command. It is recommended
to run this command when the system is under 50% CPU utilization.

== Parameters

``-node {<nodename>|local}`` - Node on Which the Profiler Should Run
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the node from which PUT/GET/DELETE operations are
sent.

``-object-store-name <text>`` - Object Store Configuration Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the object store configuration that describes the
object store. The object store configuration has information about the
object store server name, port, access credentials, and provider type.

``[-object-prefix <text>]`` - Prefix Added to Each Object [.advaced-
priv]#(privilege: advanced)#::
 This parameter specifies the prefix attached to each object. There is a
prefix length restriction of 80 characters. In addition to this
restriction, follow any specific prefix length or character restrictions
that are imposed by the cloud store they plan to run this command. Refer
to the respective cloud store documentation for details.

== Examples

The following example starts the object store profiler :
```

cluster1::>storage aggregate object-store profiler start -object-store-name my-store -node my-node

```
[[ID3ea5c903d5d615d42be5980dabe77c6b]]
```

```
= storage aggregate object-store put-rate-limit modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the maximum, per node, FabricPool put rate

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate object-store put-rate-limit modify` command enables
the modification of values that can be used to limit the FabricPool put
rate or to allow it to run more aggressively. Use the
xref:{relative_path}storage-aggregate-object-store-put-rate-limit-
show.html[storage aggregate object-store put-rate-limit show] command to
display the current maximum put rate, if any, in effect.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Node name.

``[-default {true|false}]`` - Set Default Values [.advaced-
priv]#(privilege: advanced)#::
Setting this field to true will reset all unset fields to the the system
default.

``[-put-rate-bytes-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Limit Max (B/s)
Allowed [.advaced-priv]#(privilege: advanced)#::
This field is the maximum put rate in bytes per second. The field can be
used to reduce the put rate to an object store.

``[-put-latency-limit <integer>]`` - Latency (ms) At Which Parallel Puts
Are Reduced [.advaced-priv]#(privilege: advanced)#::
This field is the maximum put latency, in milliseconds, at which point the
number of parallel put operations are decreased. The field can be used to
reduce the put rate to an object store based on average latency. Setting
this field to zero will disable the limit.

== Related Links

* xref:{relative_path}storage-aggregate-object-store-put-rate-limit-
```

```
show.html[storage aggregate object-store put-rate-limit show]




[[IDfad7f876c1ff7dd444ab096be3eb1133]]
= storage aggregate object-store put-rate-limit show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the per node FabricPool Put Rate Limit

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage aggregate object-store put-rate-limit show` command displays
the maximum put rate for FabricPool.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Node name.

``[-put-rate-bytes-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Put Rate Limit
(B/s) [.advaced-priv]#(privilege: advanced)#::
This field represents an upper boundary on the allowed put rate. The put
rate will not be allowed to exceed this value but can be throttled to
values less than it to accommodate client operations and a high average
system CPU. Some platforms are intentionally configured with a lower
maximum put rate in an effort to reduce the impact of the FabricPool put
subsystem on the overall system health. A limit in effect, lower than the
default, may keep cold data in the performance tier for a long period of
```

```
time and may require an increase in space requirements. A value of -1
indicates no upper bound is in effect.

``[-put-rate {<integer>[KB|MB|GB|TB|PB]}]`` - Estimated Put Rate (B/s)
[.advaced-priv]#(privilege: advanced)#::
The put-rate displays the maximum computed bytes per second the put
subsystem assumes is possible. This put rate is measured in real-time and
is in relation to the node for all attached object stores which have put
traffic to them.

``[-max-put-rate-in-effect {<integer>[KB|MB|GB|TB|PB]}]`` - Current
Maximum Put Rate In Effect (B/s) [.advaced-priv]#(privilege: advanced)#::
This field displays the maximum allowed put rate. If it is less than the
max-throughput then it indicates that throttling is in effect. The
throttling is the outcome of either staying at put-rates-byte-limit or an
effort to accommodate client operations on a system with high CPU.

``[-max-throughput {<integer>[KB|MB|GB|TB|PB]}]`` - Estimated Maximum
Throughput (B/s) [.advaced-priv]#(privilege: advanced)#::
This field displays the maximum put throughput rate in bytes per second
that the put subsystem will throttle down from or up to. It's the greater
of the put-rate and the node default value.

``[-put-latency <integer>]`` - Average PUT Op Latency (ms) [.advaced-
priv]#(privilege: advanced)#::
This field displays the average put latency in milliseconds over the last
several seconds to all object stores to which there is put traffic.

``[-put-latency-limit <integer>]`` - Latency at which Parallel Puts are
Reduced (ms) [.advaced-priv]#(privilege: advanced)#::
This field displays the put latency threshold in milliseconds at which
point the number of parallel put operations to all object stores is
decreased. The average put latency is displayed in the put-latency field.




[[IDa2cb6c670c7296cf4298ee360f670a5d]]
= storage aggregate plex delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a plex
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate plex delete` command deletes the specified plex.
The aggregate specified with then `-aggregate` will be unmirrored and
contain the remaining plex. The disks in the deleted plex become spare
disks.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
Name of an existing aggregate which contains the plex specified with the
`-plex` parameter.

``-plex <text>`` - Plex::
Name of a plex which belongs to the aggregate specified with the `-
aggregate` parameter.

== Examples

The following example deletes plex0 of aggregate aggr1:
```

cluster1::> storage aggregate plex delete -aggregate aggr1 -plex plex0

```
[[IDbfb943497b2363cb357e1920dbdb960a]]
= storage aggregate plex offline
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Offline a plex

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate plex offline` command takes the specified plex
offline. The aggregate specified with the `-aggregate` parameter must be a
mirrored aggregate and both plexes must be online. Prior to taking a plex
offline, the system will flush all internally-buffered data associated
with the plex and create a snapshot that is written out to both plexes.
The snapshot allows for efficient resynchronization when the plex is
subsequently brought back online.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
Name of an existing aggregate which contains the plex specified with the
`-plex` parameter.

``-plex <text>`` - Plex::
Name of a plex which belongs to the aggregate specified with the `-
aggregate` parameter.

== Examples

The following example takes plex0 of aggregate aggr1 offline:
```

cluster1::> storage aggregate plex offline -aggregate aggr1 -plex plex0

```
[[ID1a5ae25fc8f9c4c95868a5ee0f90d54c]]
= storage aggregate plex online
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Online a plex

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate plex online` command brings the specified plex
online. The aggregate specified with the `-aggregate` parameter must be an
online mirrored aggregate. The system will initiate resynchronization of
the plex as part of online processing.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
Name of an existing aggregate which contains the plex specified with the
`-plex` parameter.

``-plex <text>`` - Plex::
Name of a plex which belongs to the aggregate specified with the `-
aggregate` parameter.

== Examples

The following example brings plex0 of aggregate aggr1 online:
```

cluster1::> storage aggregate plex online -aggregate aggr1 -plex plex0

```
[[ID7bad4c95af4fef0254198d5df44253f7]]
= storage aggregate plex show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show plex details
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate plex show` command displays information for the
specified plex. By default, the command displays the following information
about all plexes:

* Aggregate Name
* Plex Name
* Is Online
* Is Resyncing
* Resyncing Percentage
* Plex Status

To display detailed information about a single plex, use the `-aggregate`
and `-plex` parameter.

== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
Name of an existing aggregate which contains the plex specified with the
`-plex` parameter.

``[-plex <text>]`` - Plex Name::
Name of a plex which belongs to the aggregate specified with the `-
aggregate` parameter.

``[-status <text>]`` - Status::
Displays plex status. Possible values are:
+
  * ``_normal_``
  * ``_failed_``
  * ``_empty_``
  * ``_invalid_``
  * ``_uninitialized_``
```

```
  * ``_failed assimilation_``
  * ``_limbo_``
  * ``_active_``
  * ``_inactive_``
  * ``_resyncing_``

+
These values may appear by themselves or in combination separated by
commas, for example, "``_normal_`` ,``_active_`` ".

``[-is-online {true|false}]`` - Is Online::
Selects the plexes that match this parameter value.

``[-in-progress {true|false}]`` - Resync is in Progress::
Selects the plexes that match this parameter value.

``[-resyncing-percent <percent>]`` - Resyncing Percentage::
Selects the plexes that match this parameter value.

``[-resync-level <integer>]`` - Resync Level::
Selects the plexes that match this parameter value.

``[-pool <integer>]`` - Pool::
Selects the plexes that match this parameter value.

== Examples

The following example displays information about all the plexes for all
the aggregates:
```

cluster1::> storage aggregate plex show
Is Is Resyncing
Aggregate Plex Online Resyncing Percent Status
--------- --------- ------- ---------- --------- ----------
aggr0 plex0 true false - normal,active
aggr1 plex0 true false - normal,active
aggr1 plex1 true false - normal,active
aggr2 plex0 true false - normal,active
aggr2 plex2 true false - normal,active
5 entries were displayed.

```
The following example displays information about plex1 of aggregate aggr1:
```

cluster1::> storage aggregate plex show -aggregate aggr1 -plex plex1
Aggregate: aggr1
Plex Name: plex1
Status: normal,active

Is Online: true
Resync is in Progress: false
Resyncing Percentage: -
Resync Level: -
Pool: 1

```
[[ID5e36aceb15526e44554aee1faef61c62]]
= storage aggregate reallocation quiesce
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Quiesce reallocate job on aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Temporarily stops any reallocation jobs that are in progress. When you use
this command, the persistent state is saved. You can use the
xref:{relative_path}storage-aggregate-reallocation-restart.html[storage
aggregate reallocation restart] command to restart a job that is quiesced.

There is no limit to how long a job can remain in the quiesced (paused)
state.

== Parameters

``-aggregate <aggregate name>`` - Aggregate Name::
Specifies the aggregate on which you want to temporarily pause the job.

== Examples
```

cluster1::> storage aggregate reallocation quiesce
-aggregate aggr0

```
Temporarily stops (pauses) any reallocation job running on aggregate
aggr0.

== Related Links

* xref:{relative_path}storage-aggregate-reallocation-restart.html[storage
aggregate reallocation restart]




[[IDb405d4eb6794d7091899e69744804d99]]
= storage aggregate reallocation restart
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restart reallocate job on aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Starts a reallocation job. Use this command to restart a quiesced
(temporarily stopped) job or a scheduled scan that is idle for the
aggregate.

== Parameters

``-aggregate <aggregate name>`` - Aggregate Name::
Specifies the aggregate on which you want to restart reallocation scans.

``[-i, -ignore-checkpoint <true>]`` - Ignore Checkpoint::
Restarts the job at the beginning when set to true. If you use this
command without specifying this parameter, its effective value is false
and the job starts the scan at the point where it was stopped. If you
specify this parameter without a value, it is set to true and the scan
restarts at the beginning.

== Examples
```

cluster1::> storage aggregate reallocation restart
-aggregate aggr0 -ignore-checkpoint true

```
Restarts reallocation job on aggregate aggr0 from the beginning.




[[IDaa9766b71a0b775773aea7a1881656eb]]
= storage aggregate reallocation schedule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify schedule of reallocate job on aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Schedules a reallocation scan for an existing reallocation job. If the
reallocation job does not exist, use the xref:{relative_path}storage-
aggregate-reallocation-start.html[storage aggregate reallocation start]
command to define a reallocation job.

You can delete an existing reallocation scan schedule. However, if you do
this, the job's scan interval reverts to the schedule that was defined for
it when the job was created with the xref:{relative_path}storage-
aggregate-reallocation-start.html[storage aggregate reallocation start]
command.

== Parameters

``-aggregate <aggregate name>`` - Aggregate Name::
Specifies the aggregate on which you want to schedule reallocation jobs.

``[-d, -del <true>]`` - Delete::
Deletes an existing reallocation schedule when set to true. If you use
this command without specifying this parameter, its effective value is
false and the reallocation schedule is not deleted. If you specify this
parameter without a value, it is set to true and the reallocation schedule
is deleted.

``[-s, -cron <text>]`` - Cron Schedule::
Specifies the schedule with the following four fields in sequence. Use a
space between field values. Enclose the values in double quotes.
```

```
+
  * minute is a value from 0 to 59.
  * hour is a value from 0 (midnight) to 23 (11:00 p.m.).
  * day of week is a value from 0 (Sunday) to 6 (Saturday).
  * day of month is a value from 1 to 31.
+
NOTE: If you specify 31 as the value for the day of month, reallocation
scans will not run in any months with fewer than 31 days.
+




+
Use an asterisk "*" as a wildcard to indicate every value for that field.
For example, an * in the day of month field means every day of the month.
You cannot use the wildcard in the minute field.
+
You can enter a number, a range, or a comma-separated list of values for a
field.

== Examples
```

cluster1::> storage aggregate reallocation schedule -aggregate aggr0 -cron "0 23 6 *"

```
Schedules a reallocation job to run at 11:00 p.m. every Saturday on aggr0.

== Related Links

* xref:{relative_path}storage-aggregate-reallocation-start.html[storage
aggregate reallocation start]




[[ID18131657aa9ef6bd9b47421d23e6e8c2]]
= storage aggregate reallocation show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show reallocate job status for improving free space layout

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

Displays the status of a reallocation scan, including the state, schedule,
aggregate and scan id. If you do not specify the `id` for a particular
reallocation scan, the command displays information about all the existing
reallocation scans.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the value of relevant field that you specify for the reallocation
scans that are present.

| ``[-v  ]``::
Specify this parameter to display the output in a verbose format.

| ``[-instance  ]`` }::
Displays information about reallocation scans on aggregates in a list
format.

``[-id <integer>]`` - Job ID::
Specify this parameter to display the reallocation scan that matches the
reallocation job ID that you specify.

``[-aggregate <aggregate name>]`` - Aggregate Name::
Specify this parameter to display the reallocation scan that matches the
aggregate that you specify.

``[-description <text>]`` - Job Description::
Specify this parameter to display reallocation scans that match the text
description that you specify.

``[-state
{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Re
schedule|Error|Quit|Dead|Unknown|Restart|Dormant}]`` - Job State::
Specify this parameter to display reallocation jobs that match the state
that you specify.

``[-progress <text>]`` - Execution Progress::
Specify this parameter to list the running reallocation jobs whose
progress indicator matches the text that you provide. For example, if you
specify "Starting ..." as the text string for the progress option, then
the system lists all the jobs that are starting.

``[-schedule <job_schedule>]`` - Schedule Name::
Specify this parameter to display reallocation scans that match the
```

```
schedule name that you specify. If you want a list of all job schedules,
use the xref:{relative_path}job-schedule-show.html[job schedule show]
command.

``[-global-status <text>]`` - Global State of Scans::
Specify this parameter to indicate if reallocation scans are on or off
globally. You must type either of the following text strings:
+
  * "Reallocation scans are on"
  * "Reallocation scans are off"



== Examples
```

cluster1::> storage aggregate reallocation show
Job ID Aggregate Schedule State
------ --------- -------- -----
23 aggr0 reallocate_0 23 * 6 Queued

```
Displays the job ID, aggregate, schedule, and state for the reallocation
scans.

== Related Links

* xref:{relative_path}job-schedule-show.html[job schedule show]




[[ID39b5d0b5f9ee17e09d4db074c5c3bfed]]
= storage aggregate reallocation start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start reallocate job on aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Begins a reallocation scan on a specified aggregate.
```

```
Before performing a reallocation scan, the reallocation job normally
performs a check of the current layout optimization. If the current layout
optimization is less than the threshold, then the system does not perform
a reallocation on the aggregate.

You can define the reallocation scan job so that it runs at a specific
interval, or you can use the xref:{relative_path}storage-aggregate-
reallocation-schedule.html[storage aggregate reallocation schedule]
command to schedule reallocation jobs.

== Parameters

``-aggregate <aggregate name>`` - Aggregate Name::
Specify this parameter to specify the target aggregate on which to start a
reallocation scan.

{ ``[-i, -interval <text>]`` - Interval Schedule::
Specified the schedule in a single string with four fields:
+
  * minute is a value from 0 to 59.
  * hour is a value from 0 (midnight) to 23 (11:00 p.m.).
  * day of month is a value from 1 to 31.
+
NOTE: If you specify 31 as the value for the day of the month,
reallocation scans will not run in any of the months with fewer than 31
days.
+

  * day of the week is a value from 0 (Sunday) to 6 (Saturday).



+
Use an asterisk "*" as a wildcard to indicate every value for that field.
For example, an * in the day of month field means every day of the month.
You cannot use the wildcard in the minute field.
+
You can enter a number, a range, or a comma-separated list of values for a
field.

| ``[-o, -once <true>]`` - Once }::
Specifies that the job runs once and then is automatically removed from
the system when set to true. If you use this command without specifying
this parameter, its effective value is false and the reallocation scan
runs as scheduled. If you enter this parameter without a value, it is set
to true and a reallocation scan runs once.
```

```
== Examples
```

cluster1::> storage aggregate reallocation start -aggregate aggr0 -interval "0 23 * 6"

```
Starts a reallocation job on aggregate aggr0 at 11:00 p.m. every Saturday.

== Related Links

* xref:{relative_path}storage-aggregate-reallocation-schedule.html[storage
aggregate reallocation schedule]




[[ID2073e8c6f09eaed0f1bde86a20adce74]]
= storage aggregate reallocation stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop reallocate job on aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Stops and deletes any reallocation scan running on the specified
aggregate. This command stops and deletes in-progress, scheduled, and
quiesced scans.

== Parameters

``-aggregate <aggregate name>`` - Aggregate Name::
Specify this parameter to specify the target aggregate on which to stop
and delete a reallocation scan.

== Examples
```

cluster1::> storage aggregate reallocation stop -aggregate aggr0

```
Stops and deletes the reallocation scan on aggregate aggr0.
```

```
[[IDce855faba9c28c5ec02c9605aac42499]]
= storage aggregate relocation show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display relocation status of an aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate relocation show` command displays status of
aggregates which were relocated in the last instance of relocation
operation.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>, ...` parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-node {<nodename>|local}]`` - Node::
Selects aggregates from the specified source node.

``[-aggregate <text>]`` - Aggregate Name::
Selects the aggregates that match this parameter value.

``[-relocation-status <text>]`` - Aggregates Relocation Status::
Selects the aggregates whose relocation status matches this parameter
value.

``[-destination <text>]`` - Destination for Relocation::
Selects the aggregates that are designated for relocation on the specified
destination node.

== Examples
```

```
The following example displays the relocation status of aggregates on all
nodes in the cluster:
```

cluster1::> storage aggregate relocation show
Source Aggregate Destination Relocation Status
--- --- --- ---
node0
- - Not attempted yet
node1
aggr1 node0 Done
aggr2 node0 In progress
aggr3 node0 Not attempted yet
4 entries were displayed.

```
[[ID769f5b03d977928387b46b508d3af86f]]
= storage aggregate relocation start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Relocate aggregates to the specified destination

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate relocation start` command initiates the relocation
of the aggregates from one node to the partner node in a high-
availablility (HA) pair.

== Parameters

``-node {<nodename>|local}`` - Name of the Node that currently owns the
aggregate::
This specifies the source node where the aggregates to be relocated
reside.

``-destination {<nodename>|local}`` - Destination node::
This specifies the destination node where aggregates are to be relocated.

``-aggregate-list <aggregate name>,...`` - List of Aggregates to be
relocated::
This specifies the list of aggregate names to be relocated from source
```

```
node to destination node.

``[-override-vetoes {true|false}]`` - Override Vetoes::
This specifies whether to override the veto checks for relocation
operation. Initiating aggregate relocation with vetoes overridden will
result in relocation proceeding even if the node detects outstanding
issues that would make aggregate relocation dangerous or disruptive. The
default value is false.

``[-relocate-to-higher-version {true|false}]`` - Relocate To Higher
Version::
This specifies if the aggregates are to be relocated to a node which is
running on a higher version of ONTAP than the source node. If an aggregate
is relocated to this destination then that aggregate cannot be relocated
back to the source node till the source is also upgraded to the same or
higher ONTAP version. This option is not required if the destination node
is running on higher minor version, but the same major version. The
default value is false.

``[-override-destination-checks {true|false}]`` - Override Destination
Checks::
This specifies if the relocation operation should override the check done
on destination node. This option could be used to force a relocation of
aggregates even if the destination has outstanding issues. Note that this
could make the relocation dangerous or disruptive. The default value is
false.

``[-ndo-controller-upgrade {true|false}]`` - Relocate Aggregates for NDO
Controller Upgrade [.advaced-priv]#(privilege: advanced)#::
This specifies if the relocation operation is being done as a part of non-
disruptive controller upgrade process. Aggregate relocation will not
change the home ownerships of the aggregates while relocating as part of
controller upgrade. The default value is false.

== Examples

The following example relocates aggregates name aggr1 and aggr2 from
source node node0 to destination node node1:
```

cluster1::> storage aggregate relocation start -node node0 -destination node1 -aggregate-list aggr1, aggr2

```
[[ID12b494c400e1af44ff84a6523c0e1318]]
= storage aggregate resynchronization modify
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify aggregate resynchronization priorities

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate resynchronization modify` command can be used to
modify the resynchronization priority of an aggregate.

When the number of aggregates pending resynchronization is higher than the
maximum number of concurrent resynchronization operations allowed on a
node, the aggregates get resynchronized in the order of their "`resync-
priority` " values.

For example, let the `max-concurrent-resync` under the `storage aggregate
resynchronization options` directory for a node be set to two. If there
are three aggregates waiting to be resynchronized, where their respective
`resync-priority` values are `high` , `medium` , and `low` , then the
third aggregate is not allowed to start resynchronization until one of the
first two aggregates has completed resynchronizing.

== Parameters

``-aggregate <aggregate name>`` - Aggregate::
This parameter specifies the aggregate that is to be modified.

``[-resync-priority {high(fixed)|high|medium|low}]`` - Resynchronization
Priority::
This parameter specifies the new resynchronization priority value for the
specified aggregate. This field cannot be modified for unmirrored or ONTAP
system aggregates.
+
Possible values for this parameter are:
+

+
  * high: Mirrored data aggregates with this priority value start
resynchronization first.
  * medium: Mirrored data aggregates with this priority value start
resynchronization after all the system aggregates and data aggregates with
'high' priority value have started resynchronization.
  * low: Mirrored data aggregates with this priority value start
resynchronization only after all the other aggregates have started
```

```
resynchronization.



== Examples

The following example changes the `resync-priority` of a specified
aggregate to `medium` :
```

cluster1::> storage aggregate resynchronization modify -aggregate aggr1 -resync-priority medium

```
[[ID5f551a6a5d1bd9677499352fbdce261e]]
= storage aggregate resynchronization show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregate resynchronization priorities

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate resynchronization show` command displays the
relative resynchronization priority for each aggregate in the cluster.
When a particular node restricts how many resync operations can be active
concurrently, these priorities are used to prioritize the operations. The
maximum concurrent resync operations for a node is displayed in the
xref:{relative_path}storage-aggregate-resynchronization-options-
show.html[storage aggregate resynchronization options show] command. If no
parameters are specified, the command displays the following information
about all the aggregates in the cluster:

* Aggregate name
* Node that owns the aggregate
* Resync priority for the aggregate



== 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.

``[-aggregate <aggregate name>]`` - Aggregate::
If this parameter is specified, the command displays the resynchronization
priority only for the specified aggregate.

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays the resynchronization
priority only for the aggregates owned by the specified node.

``[-resync-priority {high(fixed)|high|medium|low}]`` - Resynchronization
Priority::
If this parameter is specified, the command displays only the
resynchronization priority that matches the specified value. Possible
values for this parameter are:
+
  * high(fixed): This value is reserved for ONTAP system aggregates, which
cannot have any other value for this field. These aggregates always start
their resynchronization operation at the first available opportunity. This
value cannot be assigned to a data aggregate.
  * high: Mirrored data aggregates with this priority value start
resynchronization first.
  * medium: Mirrored data aggregates with this priority value start
resynchronization after all the system aggregates and data aggregates with
'high' priority value have started resynchronization.
  * low: Mirrored data aggregates with this priority value start
resynchronization only after all the other aggregates have started
resynchronization.

+
When the number of aggregates waiting for resynchronization is higher than
the maximum number of resynchronization operations allowed on a node, then
the resync-priority field is used to determine which aggregate starts
resynchronization first. This field is not set for unmirrored aggregates.

== Examples

The following command displays the resynchronization priorities for all
the aggregates in the cluster:
```

cluster1::> storage aggregate resynchronization show

Aggregate Node Resync Priority
--------- --------------- ---------------
aggr0_n1 cluster1-01 high(fixed)
aggr0_n2 cluster1-02 high(fixed)
aggr1 cluster1-01 low
aggr2 cluster1-01 high
aggr3 cluster1-01 medium
4 entries were displayed.

```
== Related Links

* xref:{relative_path}storage-aggregate-resynchronization-options-
show.html[storage aggregate resynchronization options show]




[[ID03703283da7a682f51720b1bd5ea19f9]]
= storage aggregate resynchronization options modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify node specific aggregate resynchronization options

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate resynchronization options modify` command can be
used to modify the options that govern the resynchronization of aggregates
on a given cluster node.

Modifying the `max-concurrent-resyncs` option changes the number of
aggregates that are allowed to resynchronize concurrently. When the number
of aggregates waiting for resynchronization is higher than this value, the
aggregates are resynchronized in the order of their "`resync-priority` ".
This value can be modified using the xref:{relative_path}storage-
aggregate-resynchronization-modify.html[storage aggregate
resynchronization modify] command while specifying the `-resync-priority`
parameter.

== Parameters
```

```
``-node {<nodename>|local}`` - Node::
This parameter specifies the node for which the option is to be modified.

``[-max-concurrent-resync <integer>]`` - Maximum Concurrent
Resynchronizing Aggregates::
This parameter specifies the new value for the maximum number of
concurrent resync operations allowed on a node. This option must be
specified along with the `-node` parameter. When a node has active resync
operations, setting this parameter to a value that is lower than the
number of currently resyncing aggregates will trigger a user confirmation.

== Examples

The following example changes the maximum concurrent resync operations for
the specified node to `ten` :
```

cluster1::> storage aggregate resynchronization options modify -node node1 -max-concurrent-resyncs 10

```
== Related Links

* xref:{relative_path}storage-aggregate-resynchronization-
modify.html[storage aggregate resynchronization modify]




[[ID53f45f80c68d8a44853061d5c40bca41]]
= storage aggregate resynchronization options show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display node specific aggregate resynchronization options

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage aggregate resynchronization options show` command displays
all the options that govern the resynchronization of aggregates on a given
cluster node. If no parameters are specified, the command displays the
following information about all nodes:
```

```
* Node for which the information is being displayed
* Maximum number of concurrent resynchronizing aggregates allowed



== 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>|local}]`` - Node::
If this parameter is specified, the command displays resynchronization
options only for the specified node.

``[-max-concurrent-resync <integer>]`` - Maximum Concurrent
Resynchronizing Aggregates::
If this parameter is specified, the command displays only the
resynchronization option that matches the specified value.

== Examples

The following example displays the maximum number of concurrent resyncs
allowed for each node in the cluster:
```

cluster1::> storage aggregate resynchronization options show
Node Maximum Concurrent Resynchronizing Aggregates
--- ---
cluster1-01 15
cluster1-02 4
2 entries were displayed.

```
The following example displays the maximum number of concurrent resyncs
allowed for a specified node:
```

cluster1::> storage aggregate resynchronization options show -node node1
Node Maximum Concurrent Resynchronizing Aggregates
--- ---
cluster1-01 15

```
The following example displays all the nodes that allow more than five
concurrent resync operations:
```

cluster1::> storage aggregate resynchronization options show -max-concurrent-resyncs >5
Node Maximum Concurrent Resynchronizing Aggregates
---------- ----------
cluster1-01 15

```
:leveloffset: -1


= storage array commands

:leveloffset: +1


[[ID9ef50bce8db90a7a88a39e168e511eef]]
= storage array modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Make changes to an array's profile.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array modify` command lets the user change several array
parameters.

== Parameters

``-name <text>`` - Name::
Storage array name, either generated by ONTAP or assigned by the user.

``[-prefix <text>]`` - Prefix::
Abbreviation for the named array.

``[-vendor <text>]`` - Vendor::
Array manufacturer.
```

```
``[-model <text>]`` - Model::
Array model number.

``[-options <text>]`` - options::
Vendor specific array settings.

``[-max-queue-depth <integer>]`` - Target Port Queue Depth [.advaced-
priv]#(privilege: advanced)#::
The target port queue depth for all target ports on this array.

``[-lun-queue-depth <integer>]`` - LUN Queue Depth [.advaced-
priv]#(privilege: advanced)#::
The queue depth assigned to array LUNs from this array.

{ ``[-is-upgrade-pending {true|false}]`` - Upgrade Pending [.advaced-
priv]#(privilege: advanced)#::
Set this parameter to ``_true_`` if the array requires additional ONTAP
resilience for a pending firmware upgrade. Keep this parameter ``_false_``
during normal array operation. This value can not be set to ``_true_`` if
-path-faliover-time is greater than zero.

``[-path-failover-time <integer>]`` - Path Failover Time (sec)::
The time delay (in secs) before switching the I/O path when the path is
deleted. The maximum time delay is 30 sec. The default is 0. This value
can not be greater than zero if -is-upgrade-pending is ``_true_`` .

``[-all-path-fail-delay <integer>]`` - Extend All Path Failure Event
(secs)::
Use this parameter to increase the delay before ONTAP declares an "all
path failure" event for an array. Delaying the "all path failure" event
allows ONTAP to suspend I/O operations for a longer period of time before
declaring a data access disruption, allowing for I/O operations to resume
if any path comes back online within the specified duration. A valid delay
is any value between 30 and 90 seconds. A value of 0 will reset the delay,
resulting in default actions being taken whenever an "all path failure"
event is detected.

== Examples

This command changes the model to FastT.
```

cluster1::> storage array modify -name IBM_1722_1 -model FastT

```
[[ID78bb752a963c1302f1974be6562a7888]]
= storage array remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a storage array record from the array profile database.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array remove` command discards array profile records for a
particular storage array from the cluster database. The command fails if a
storage array is still connected to the cluster. Use the
xref:{relative_path}storage-array-config-show.html[storage array config
show] command to view the array connectivity status. The array target port
can be removed using the xref:{relative_path}storage-array-port-
remove.html[storage array port remove] command.

== Parameters

``-name <text>`` - Name::
Name of the storage array you want to remove from the database.

== Examples
```

cluster1::> storage array remove IBM_1722_1

```
== Related Links

* xref:{relative_path}storage-array-config-show.html[storage array config
show]
* xref:{relative_path}storage-array-port-remove.html[storage array port
remove]




[[IDb3b7d1f79a83d9a3c477a63b98aa5132]]
= storage array rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change the name of a storage array in the array profile database.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array rename` command permits substitution of the array
profile name which ONTAP assigned during device discovery. By default, the
name that ONTAP assigned to the storage array during discovery is shown in
ONTAP displays and command output.

== Parameters

``-name <text>`` - Name::
Storage array name either generated by ONTAP or assigned by the user.

``-new-name <text>`` - The new name to assign to this array profile. (28
chars max)::
New name to assign to the storage array.

== Examples
```

cluster1::> storage array rename -name HITACHI_DF600F_1 -new-name MyArray

```
[[ID8553816d7583ce2b55ac24c78e290fe8]]
= storage array show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about SAN-attached storage arrays.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array show` command displays information about arrays visible
to the cluster. If no parameters are specified, the command displays the
following information about all storage arrays:

* Prefix
* Name
* Vendor
* Model
* Options

 To display detailed information about a single array, use the `-name`
parameter. The detailed view  adds the following  information:

* Serial Number
* Optimization Policy
* Affinity
* Errors
* Path Failover Time
* Extend All Path Failure Event



== 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.

``[-name <text>]`` - Name::
```

```
Selects the arrays that match this parameter value.

``[-prefix <text>]`` - Prefix::
Abbreviation for the named array.

``[-vendor <text>]`` - Vendor::
Array manufacturer.

``[-model <text>]`` - Model::
Array model number.

``[-options <text>]`` - options::
Vendor specific array settings.

``[-serial-number <text>]`` - Serial Number::
Array product identifier.

``[-max-queue-depth <integer>]`` - Target Port Queue Depth [.advaced-
priv]#(privilege: advanced)#::
Selects the arrays that match this parameter value.

``[-lun-queue-depth <integer>]`` - LUN Queue Depth [.advaced-
priv]#(privilege: advanced)#::
Selects the arrays that match this parameter value.

``[-optimization-policy
{iALUA|eALUA|symmetric|proprietary|mixed|unknown}]`` - Optimization
Policy::
Selects the arrays that match this parameter value.

``[-affinity {none|aaa|ap|mixed|unknown}]`` - Affinity::
Selects the arrays that match this parameter value.

``[-error-text <text>,...]`` - Error Text::
Selects the arrays that match this parameter value.

``[-is-upgrade-pending {true|false}]`` - Upgrade Pending [.advaced-
priv]#(privilege: advanced)#::
Selects the arrays that match this parameter value.

``[-path-failover-time <integer>]`` - Path Failover Time (sec)::
Use this parameter to list arrays that have path failover time set to the
value you specify.

``[-all-path-fail-delay <integer>]`` - Extend All Path Failure Event
(secs)::
Use this parameter to list arrays that have the all path failure event
```

```
delay set to the value you specify.

== Examples

The following example displays information about all arrays.
```

cluster1::> storage array show
Prefix Name Vendor Model Options
------ ---------- -------- ---------- ----------
HITACHI_DF600F_1 HITACHI DF600F
IBM_1722_1 IBM 1722
2 entries were displayed.

```
The following example displays detailed information about a specific
array:
```

cluster1::> storage array show -name HITACHI_DF600F_1
Name: HITACHI_DF600F_1
Prefix: abc
Vendor: HITACHI
Model: DF600F
options:
Serial Number: 4291000000000000
Optimization Policy: iALUA
Affinity: aaa
Error Text:
Path Failover Timeout (sec): 30
Extend All Path Failure Event (secs): 50

```
[[ID08e91900bec76f039af3239de4fb37b8]]
= storage array config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display connectivity to back-end storage arrays.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array config show` command displays information about how the
storage arrays connect to the cluster, LUN groups, number of LUNS, and
more. Use this command to validate the configuration and to assist in
```

```
troubleshooting.

== 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.

| ``[-switch  ]``::
If you specify this parameter, switch port information is shown.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Controller Name::
Selects the arrays that match this parameter value.

``[-group <integer>]`` - LUN Group::
Selects the arrays that match this parameter value. A LUN group is a set
of LUNs that shares the same path set.

``[-target-wwpn <text>]`` - Array Target Ports::
Selects the arrays that match this parameter value (the World Wide Port
Name of a storage array port).

``[-initiator <text>]`` - Initiator::
Selects the arrays that match this parameter value (the host bus adapter
that the clustered node uses to connect to storage arrays).

``[-array-name <array name>]`` - Array Name::
Selects the arrays that match this parameter value.

``[-target-side-switch-port <text>]`` - Target Side Switch Port::
Selects the arrays that match this parameter value.

``[-initiator-side-switch-port <text>]`` - Initiator Side Switch Port::
Selects the arrays that match this parameter value.

``[-lun-count <integer>]`` - Number of array LUNs::
Selects the arrays that match this parameter value.

``[-ownership {all|assigned|unassigned}]`` - Ownership::
Selects the arrays that match this parameter value.

== Examples
```

```
cluster1::> storage array config show
LUN LUN
Node Group Count Array Name Array Target Port Initiator
------------ ----- ----- --------------------------- ---------------------- ---------
vnv3070f19a 0 20 DGC_RAID5_1 5006016030229f13 0d
5006016130229f13 0c
5006016830229f13 0b
5006016930229f13 0a
1 21 HITACHI_OPEN_1 50060e80034fe704 0c
0d
50060e80034fe714 0a
0b
50060e80034fe715 0b
50060e80034fe716 0c
0d
2 8 EMC_SYMMETRIX_1 50060482cb1bce1d 0a
0b
5006048acb1bce0c 0c
0d
3 10 IBM_UniversalXport_1 202600a0b8322d10 0c
0d
204700a0b8322d10 0a
0b
vnv3070f19b 0 20 DGC_RAID5_1 5006016030229f13 0d
5006016130229f13 0c
5006016830229f13 0b
5006016930229f13 0a
1 21 HITACHI_OPEN_1 50060e80034fe704 0c
0d
50060e80034fe714 0a
0b
50060e80034fe715 0b
50060e80034fe716 0c
0d
2 8 EMC_SYMMETRIX_1 50060482cb1bce1d 0a
0b
5006048acb1bce0c 0c
0d
3 10 IBM_UniversalXport_1 202600a0b8322d10 0c
0d
204700a0b8322d10 0a
0b
38 entries were displayed.

Warning: Configuration errors were detected. Use 'storage errors show' for detailed information.
```

```
[[ID4301fd2b11b4b874bc5bdcd78b534245]]
= storage array disk paths show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display a list of LUNs on the given array

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array disk paths show` command displays information about
disks and array LUNs. Where it appears in the remainder of this document,
"disk" may refer to either a disk or an array LUN. By default, the command
displays the following information about all disks:

* Disk Unique Identifier
* Controller name
* Initiator Port
* LUN ID
* Failover optimization type
* The Use State of the LUN on this path
* Target Port
* Target IQN
* TPGN
* Port speeds
* Kbytes/sec on Disk (Rolling Average)
* Number IOPS per second on disk (Rolling Average)

To display detailed information about a single disk, use the `-disk`
parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all disks, in column style output.

| ``[-switch  ]``::
Displays the switch port information for all disks, in column style
output.

| ``[-instance  ]`` }::
Displays detailed disk information. If no disk path name is specified,
this parameter displays the same detailed information for all disks as
does the -disk parameter. If a disk path name is specified, then this
parameter displays the same detailed information for the specified disks
as does the -disk parameter.

``[-uid <text>]`` - Disk Unique Identifier::
```

```
Selects the disks whose unique id matches this parameter value. A disk
unique identifier has the form:
`20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000 `

``[-disk <disk path name>]`` - Disk Name::
Displays detailed information about the specified disks.

``[-array-name <array name>]`` - Array Name::
 Selects information about the LUNs presented by the specified storage
array.

``[-diskpathnames <disk path name>,...]`` - Path-Based Disk Names::
Selects information about disks that have all of the specified path names.

``[-nodelist {<nodename>|local}]`` - Controller name::
Selects information about disks that are visible to all of the specified
nodes .

``[-initiator <text>,...]`` - Initiator Port::
Selects information about disks that are visible to the initiator
specified. Disks that are not currently in use by that initiator are
included.

``[-lun <integer>,...]`` - LUN ID::
Selects information about the specified LUNs.

``[-target-wwpn <text>,...]`` - Target Port::
Selects information about disks that are visible on target ports
identified by their World Wide Port Name.

``[-initiator-side-switch-port <text>,...]`` - Initiator Side Switch
Port::
Selects information about disks visible to an initiator that is connected
to the specified switch port.

``[-lun-path-use-state <text>,...]`` - The Use State of the LUN on this
path::
Selects information about LUNs reporting the specified in-use state.

``[-tpgn <integer>,...]`` - Target Port Group Number::
Selects information about disks that belong to the specified Target Port
Group Number.

``[-port-speed <text>,...]`` - Port Speed::
Selects information about disks served by a Host Bus Adapter that is
running at the specified port speed.
```

```
``[-lun-io-kbps <integer>,...]`` - Kbytes/sec on Disk (Rolling Average)::
Selects information about the LUNs that have reached the specified I/O
throughput.

``[-lun-iops <integer>,...]`` - Number IOPS per second on disk (Rolling
Average)::
Selects information about the LUNs that have reached the specified number
of IOPs.

``[-target-side-switch-port <text>,...]`` - Target Side Switch Port::
Selects information about disks that are visible on target ports
identified by the switch port to which they are connected.

``[-target-port-access-state <text>,...]`` - Failover optimization type::
Selects information about disks visible on target ports that have the
specified access state.

``[-initiator-io-kbps <integer>,...]`` - Kbytes of I/O per second on
Initiator (Rolling Average)::
Selects information about disks visible to an initiator that has executed
I/O at the specified throughput.

``[-initiator-iops <integer>,...]`` - Number of IOPS on Initiator (Rolling
Average)::
Selects information about disks visible to an initiator that has executed
the specified number of IOPs.

``[-target-io-kbps <integer>,...]`` - Kbytes of I/O per second to Target
(Rolling Average)::
Selects information about disks visible on target ports that have reached
the specified I/O throughput.

``[-target-iops <integer>,...]`` - Number of IOPS to Target (Rolling
Average)::
Selects information about disks visible on target ports that have
performed the specified number of IOPs.

``[-path-link-errors <integer>,...]`` - Link Error count on path::
Selects information about disks with paths that have incurred the
specified number of FC link errors.

``[-path-io-kbps <integer>,...]`` - Kbytes of I/O per second on Path
(Rolling Average)::
Selects information about disk with paths that have reached the specified
I/O throughput.
```

```
``[-path-iops <integer>,...]`` - Number of IOPS on Path (Rolling
Average)::
Selects information about disks on those paths that have reached the
specified number of IOPs.

``[-path-quality <integer>,...]`` - Percentage of weighted error
threshold::
Selects information about disks on paths that have incurred the specified
number of errors. The value displayed is a measure of the health of a path
expressed as a percentage of an error threshold. Once a path has reached
or surpassed the error threshold, another path will be selected for I/O
transfer, if there is one available.

``[-path-lun-in-use-count <integer>,...]`` - Number of LUNs in the in-use
state on this path::
Selects information about disks with paths that have the specified in-use-
count.

``[-initiator-lun-in-use-count <integer>,...]`` - Number of LUNs in the
in-use state on this initiator::
Selects information about disks with a path through an initiator that has
the specified in-use-count.

``[-target-lun-in-use-count <integer>,...]`` - Number of LUNs in the in-
use state on this target::
Selects information about disks with a path through a target port that has
the specified in-use-count.

``[-preferred-target-port {true|false}]`` - Whether or not target port
group is preferred::
Selects information about disks that match the specified parameter value
indicating whether the backing storage is ALUA (Asymmetric Logical Unit
Access) capable and has specified the array target port on this path to be
a preferred target port for I/O.

``[-vmdisk-device-id <integer>,...]`` - Virtual disk device ID::
Selects information about disks that have the specified virtual disk
device ID.

``[-host-adapter <text>]`` - Primary Path Host Adapter::
Selects information about disks that are currently using the specified
Host Bus Adapter.

``[-primary-port <text>]`` - Primary Path Disk Port::
Selects information about disks that use the specified primary port.

``[-secondary-name <disk path name>]`` - Secondary Path Name::
```

```
Selects information about disks that use the specified secondary path
name, for multipath configuration.

``[-secondary-port <text>]`` - Secondary Path Disk Port::
Selects information about disks that use the specified secondary port.

== Examples

The following example displays information about all disks:
```

cluster1::> storage array disk paths show
Disk Name: 1.0.20
UID:
5000C500:0979E09F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000
LUN Link Disk I/O
Controller Initiator ID Acc Use Target Port TPGN Speed (KB/s) IOPS
node2 3a 0 AO INU 5000c5000979e09d 80 9 Gb/S 0 0
node2 3c 0 AO RDY 5000c5000979e09e 12 9 Gb/S 0 0
node1 3a 0 AO RDY 5000c5000979e09e 12 9 Gb/S 0 0
node1 3c 0 AO INU 5000c5000979e09d 80 9 Gb/S 0 0
Disk Name: 1.0.22
UID:
5000C500:0979E3C3:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000
LUN Link Disk I/O
Controller Initiator ID Acc Use Target Port TPGN Speed (KB/s) IOPS
node2 3a 0 AO INU 5000c5000979e3c1 83 9 Gb/S 0 0
node2 3c 0 AO RDY 5000c5000979e3c2 15 9 Gb/S 0 0
node1 3a 0 AO RDY 5000c5000979e3c2 15 9 Gb/S 0 0
node1 3c 0 AO INU 5000c5000979e3c1 83 9 Gb/S 0 0
Disk Name: 1.0.19
UID:
5000C500:0979E3F3:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000
LUN Link Disk I/O
Controller Initiator ID Acc Use Target Port TPGN Speed (KB/s) IOPS
node2 3a 0 AO RDY 5000c5000979e3f1 86 9 Gb/S 0 0
node2 3c 0 AO INU 5000c5000979e3f2 18 9 Gb/S 0 0
node1 3a 0 AO INU 5000c5000979e3f2 18 9 Gb/S 0 0
node1 3c 0 AO RDY 5000c5000979e3f1 86 9 Gb/S 0 0
Disk Name: 1.0.16
UID:
5000C500:0979EBEB:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000
LUN Link Disk I/O
Controller Initiator ID Acc Use Target Port TPGN Speed (KB/s) IOPS
node2 3a 0 AO INU 5000c5000979ebe9 71 9 Gb/S 283 3
node2 3c 0 AO RDY 5000c5000979ebea 3 9 Gb/S 0 0
node1 3a 0 AO RDY 5000c5000979ebea 3 9 Gb/S 0 0
node1 3c 0 AO INU 5000c5000979ebe9 71 9 Gb/S 3 0

```
[[ID44cc9de29f3b64e9d54008d5efdbd28f]]
= storage array port modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Make changes to a target port record.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array port modify` command lets the user change array target
port parameters.

== Parameters

``-name <text>`` - Name::
Selects the array ports that match this parameter value. The storage array
name is either generated by ONTAP or assigned by the user.

``-wwnn <text>`` - WWNN::
Selects the array ports that match this parameter value.

``-wwpn <text>`` - WWPN::
Selects the array ports that match this parameter value.

``[-max-queue-depth <integer>]`` - Target Port Queue Depth::
The target port queue depth for this target port.

``[-utilization-policy {normal|defer}]`` - Utilization Policy::
The policy used in automatically adjusting the queue depth of the target
port based on its utilization.

== Examples

This command changes the maximum queue depth for this target port to 32.
```

cluster1::> storage array port modify -name HITACHI_DF600F_1 -wwnn 50060e80004291c0 -wwpn 50060e80004291c0 -max-queue-depth 32

```
[[ID44748acebc362ed365b0ffb03da096d2]]
= storage array port remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a port record from an array profile.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array port remove` command removes a port from the array
database. You might want to remove ports that are no longer connected to
the clustered node. Port information can change after hardware
replacement, rezoning, or similar configuration activities. The database
retains the records about previous ports unless you remove the
information.

== Parameters

``-name <text>`` - Name::
Selects the array ports that match this parameter value. The storage array
name is either generated by ONTAP or assigned by the user.

{ ``[-wwnn <text>]`` - WWNN::
Selects the array ports that match this parameter value.

``[-wwpn <text>]`` - WWPN::
Selects the array ports that match this parameter value.

| ``[-target-iqn <text>]`` - Target IQN::
Selects the array ports that match this parameter value.

``[-tpgt <integer>]`` - TPGT }::
Selects the array ports that match this parameter value.

== Examples

This command removes a port record from the array profiles database.
```

cluster1::> storage array port remove -name HITACHI_DF600F_1 -wwnn 50060e80004291c0 -wwpn 50060e80004291c0

```
[[IDcfc743c22c12f62d4085ca87974b061f]]
= storage array port show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about a storage array's target ports.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage array port show` command displays all the target ports known
to the cluster for a given storage array (if an array name is specified)
or for all storage arrays if no storage array name is specified. Target
ports remain in the database as part of an array profile unless you
explicitly remove them from the database.

== 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.

``[-name <text>]`` - Name::
Selects the array ports that match this parameter value. The storage array
name is either generated by ONTAP or assigned by the user.

``[-wwnn <text>]`` - WWNN::
Selects the array ports that match this parameter value.

``[-wwpn <text>]`` - WWPN::
Selects the array ports that match this parameter value.

``[-max-queue-depth <integer>]`` - Target Port Queue Depth::
Selects the array ports that match this parameter value.

``[-node {<nodename>|local}]`` - Controller Name::
```

```
Selects the array ports that match this parameter value.

``[-initiator-port <text>,...]`` - Initiator Port::
Selects the array ports that match this parameter value.

``[-average-dynamic-queue-depth <integer>,...]`` - Average Dynamic Queue
Depth [.advaced-priv]#(privilege: advanced)#::
The average value of the dynamic target port queue depth.

``[-average-latency-per-iop <integer>,...]`` - Average Latency Per IOP::
Selects the array ports that match this parameter value (average latency
per I/O performed in microseconds).

``[-average-pending <integer>,...]`` - Average Pending [.advaced-
priv]#(privilege: advanced)#::
Selects the array ports that match this parameter value (average over time
of how many commands are on the outstanding queue).

``[-average-waiting <integer>,...]`` - Average Waiting [.advaced-
priv]#(privilege: advanced)#::
Selects the array ports that match this parameter value (average over time
of how many commands are on the waiting queue).

``[-connection-type {direct|fabric|ISCSI}]`` - Connection Type::
Selects the array ports that match this parameter value (type of
connection between the controller and the back end storage).

``[-dynamic-queue-depth <integer>,...]`` - Dynamic Queue Depth [.advaced-
priv]#(privilege: advanced)#::
Current dynamic target port queue depth, the maximum number of commands
allowed outstanding.

``[-max-pending <integer>,...]`` - Max Pending [.advaced-priv]#(privilege:
advanced)#::
Selects the array ports that match this parameter value (largest number of
commands observed on the outstanding queue).

``[-max-waiting <integer>,...]`` - Max Waiting [.advaced-priv]#(privilege:
advanced)#::
Selects the array ports that match this parameter value (largest number of
commands observed on the waiting queue).

``[-path-link-errors <integer>,...]`` - Link Error count on path::
Selects the array ports that match this parameter value.

``[-percent-busy <integer>,...]`` - Percent Busy::
Selects the array ports that match this parameter value (percentage of
```

```
time I/Os are outstanding on the port).

``[-percent-waiting <integer>,...]`` - Percent Waiting::
Selects the array ports that match this parameter value (percentage of
time there are I/Os waiting on the throttle list on the target port).

``[-switch-port <text>]`` - Switch Port::
Selects the array ports that match this parameter value (for fabric
attached connections, the switch port the array target port is connected
to; N/A for direct attached).

``[-target-io-kbps <integer>,...]`` - Kbytes of I/O per second to Target
(Rolling Average)::
Selects the array ports that match this parameter value.

``[-target-iops <integer>,...]`` - Number of IOPS to Target (Rolling
Average)::
Selects the array ports that match this parameter value.

``[-target-lun-in-use-count <integer>,...]`` - Target LUN In Use Count::
Selects the array ports that match this parameter value (number of IN-USE
disks on this target port).

``[-target-port-speed <text>]`` - Target Port Speed::
Selects the array ports that match this parameter value (speed that the
target port has negotiated with its connected switch port, or initiator
port if direct attached).

``[-utilization-policy {normal|defer}]`` - Utilization Policy::
The policy used when sending I/O to an array target port when it reaches
maximum queue depth. Possible values are:
+
  * normal - This policy aggressively competes for target port resources,
in effect competing with other hosts. (default)
  * defer - This policy does not aggressively compete for target port
resources, in effect deferring to other hosts.



== Examples

The example below displays the port information for a single port.
```

cluster1::> storage array port show -wwpn 50060e80004291c0
Array Name: HITACHI_DF600F_1
WWNN: 50060e80004291c0
WWPN: 50060e80004291c0

Connection Type: fabric
Switch Port: vgbr300s89:9
Link Speed: 4 GB/s
Max Queue Depth: 1024
Utilization Policy: normal
LUN Link
Node Initiator Count IOPS KB/s %busy %waiting Errs
---------- --------- ----- ------ ------ ------- ---------- ----
vnv3070f20a 0b 2 0 0 0 0 0
vnv3070f20b 0b 2 0 0 0 0 0

```
:leveloffset: -1


= storage automated-working-set-analyzer commands

:leveloffset: +1


[[ID05fe9aee386b4b8da1433aa780d13d87]]
= storage automated-working-set-analyzer show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-node <nodename>]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the node name that the AWA instance runs on.

``[-flash-cache {true|false}]`` - Flash Cache Node-wide Modeling
[.advaced-priv]#(privilege: advanced)#::
This parameter indicates whether the AWA is modeling flash-cache.

``[-aggregate-uuid <UUID>]`` - Uuid of the Aggregate [.advaced-
priv]#(privilege: advanced)#::
This parameter indicates the aggregate uuid that the AWA instance runs on.

``[-aggregate <aggregate name>]`` - Aggregate [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the aggregate name that the AWA instance runs on.

``[-working-set-size {true|false}]`` - Working Set Size [.advaced-
priv]#(privilege: advanced)#::
This parameter indicates whether the AWA instance is configured to find
the working set size.

``[-start-time <Date>]`` - Starting Time [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the time when the AWA instance was started.

``[-total-intervals <integer>]`` - Total Interval Count [.advaced-
priv]#(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
[.advaced-priv]#(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
[.advaced-priv]#(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 [.advaced-
priv]#(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 [.advaced-
```

```
priv]#(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 [.advaced-priv]#(privilege: advanced)#::
This parameter indicates the projected Flash Pool cache usage.

``[-projected-read-hit <percent>]`` - Projected Read Hit [.advaced-
priv]#(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 [.advaced-
priv]#(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
[.advaced-priv]#(privilege: advanced)#::
This parameter indicates the interval in which the cache size effect
information is derived from.

``[-referenced-interval-time <Date>]`` - Referenced Interval Time
[.advaced-priv]#(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 [.advaced-priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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
[.advaced-priv]#(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
```

```
[[ID700b3a68721f5f3530f081ba297b9c5a]]
= storage automated-working-set-analyzer start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
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 [.advaced-priv]#(privilege: advanced)#::
This parameter indicates the node name that the AWA instance runs on.

``[-flash-cache {true|false}]`` - Flash Cache Node-wide Modeling
[.advaced-priv]#(privilege: advanced)#::
This parameter indicates whether the AWA is modeling flash-cache.

``[-aggregate <aggregate name>]`` - Aggregate [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the aggregate name that the AWA instance runs on.

``[-working-set-size {true|false}]`` - Working Set Size [.advaced-
priv]#(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

```
[[IDb8f39958296cb961796cd233450ef867]]
= storage automated-working-set-analyzer stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
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 [.advaced-priv]#(privilege: advanced)#::
This parameter indicates the node name that the AWA instance runs on.

``[-flash-cache {true|false}]`` - Flash cache node-wide modeling
[.advaced-priv]#(privilege: advanced)#::
This parameter indicates whether the AWA is modeling flash-cache.

``[-aggregate <aggregate name>]`` - Aggregate [.advaced-priv]#(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

```
[[IDfb8b00dfd0f71e406741e081935dde9c]]
= storage automated-working-set-analyzer volume show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - Node [.advaced-priv]#(privilege: advanced)#::
This parameter indicates the node name that the AWA instance runs on.

``[-flash-cache {true|false}]`` - Flash Cache Node-wide Modeling
[.advaced-priv]#(privilege: advanced)#::
This parameter indicates whether the AWA is modeling flash-cache.

``[-vol-uuid <UUID>]`` - Uuid of the Volume [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the volume uuid that this command is issued on.

``[-aggregate <aggregate name>]`` - Aggregate [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the aggregate name that the AWA instance runs on.

``[-volume <volume name>]`` - Volume [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the volume name that this command is issued on.

``[-rank <integer>]`` - Cache Benefit Rank [.advaced-priv]#(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
[.advaced-priv]#(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
[.advaced-priv]#(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 [.advaced-
priv]#(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 [.advaced-
priv]#(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 [.advaced-priv]#(privilege: advanced)#::
This parameter indicates the projected Flash Pool cache usage by this
volume.

``[-projected-read-hit <percent>]`` - Projected Read Hit [.advaced-
priv]#(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 [.advaced-
priv]#(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
[.advaced-priv]#(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

```
:leveloffset: -1


= storage disk commands

:leveloffset: +1


[[IDc1cff0d412793b4a1848af87f6c35ef7]]
= storage disk assign
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Assign ownership of a disk to a system

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk assign` command is used to assign ownership of an
unowned disk or array LUN to a specific node. You can also use this
command to change the ownership of a disk or an array LUN to another node.
You can designate disk ownership by specifying disk names, array LUN
names, wildcards, or all (for all disks or array LUNs visible to the
node). For disks, you can also set up disk ownership autoassignment. You
can also assign disks to a particular pool. You can also assign disks by
copying ownership from another disk.

== Parameters

{ ``[-disk <disk path name>]`` - Disk Path::
This specifies the disk or array LUN that is to be assigned. Disk names
take one of the following forms:
+
  * Disks are named in the form ``_<stack-id>_`` .``_<shelf>_``
.``_<bay>_``
  * Disks on multi-disk carriers are named in the form ``_<stack-id>_``
.``_<shelf>_`` .``_<bay>_`` .``_<lun>_``
```

```
  * Virtual disks are named in the form <prefix>.<number>, where prefix is
the storage array's prefix and number is a unique ascending number.

+
Disk names take one of the following forms on clusters that are not yet
fully upgraded to ONTAP 8.3:
+
  * Disks that are not attached to a switch are named in the form
``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_`` . For disks with a
LUN, the form is ``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_``
L``_<LUN>_`` . For instance, disk number 16 on host adapter 1a on a node
named node0a is named node0a:1a.16. The same disk on LUN lun0 is named
node0a:1a.16Llun0.
  * Disks that are attached to a switch are named in the form ``_<node>_``
:``_<switch_name>_`` :``_<switch_port>_`` .``_<loop_ID>_`` . For disks
with a LUN, the form is ``_<node>_`` :``_<switch_name>_``
:``_<switch_port>_`` .``_<loop_ID>_`` L``_<LUN>_`` . For instance, disk
number 08 on port 11 of switch fc1 on a node named node0a is named
node0a:fc1:11.08. The same disk on LUN lun1 is named
node0a:fc1:11.08Llun1.



+
Before the cluster is upgraded to ONTAP 8.3, the same disk can have
multiple disk names, depending on how the disk is connected. For example,
a disk known to a node named alpha as alpha:1a.19 can be known to a node
named beta as beta:0b.37. All names are listed in the output of queries
and are equally valid. To determine a disk's unique identity, run a
detailed query and look for the disk's universal unique identifier (UUID)
or serial number.
+
A subset of disks or array LUNs can be assigned using the wildcard
character (*) in the `-disk` parameter. Either the `-owner` , the `-sysid`
, or the `-copy-ownership-from` parameter must be specified with the `-
disk` parameter. Do not use the `-node` parameter with the `-disk`
parameter.

| ``-disklist <disk path name>,...`` - Disk list::
This specifies the List of disks to be assigned.

| ``-all <true>`` - Assign All Disks::
This optional parameter causes assignment of all visible unowned disks or
array LUNs to the node specified in the `-node` parameter. The `-node`
parameter must be specified with the `-all` parameter. When the `-copy-
ownership-from` parameter is specified with the `-node` parameter, it
assigns disk ownership based on the `-copy-ownership-from` parameter;
```

```
otherwise it assigns ownership of the disks based on the `-node`
parameter. Do not use the `-owner` or the `-sysid` parameter with the `-
all` parameter.

| ``[-T, -type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD |
VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Storage
Type::
This optional parameter assigns ownership of a specific type of disk or
array LUN (or a set of disks/array LUNs) to a node. The `-count` parameter
must be specified with the `-type` parameter.

``-n, -count <integer>`` - Disk Count::
This optional parameter assigns ownership of a number of disks or array
LUNs specified in the `-count` parameter, to a node.

| ``-auto <true>`` - Auto Assign }::
This optional parameter causes all visible disks eligible for
autoassignment to be immediately assigned to the node specified in the `-
node` parameter. Only unowned disks on loops or stacks owned wholly by
that system and which have the same pool information will be assigned. The
`-node` parameter must be specified with the `-auto` parameter. Do not use
the `-owner` , the `-sysid` , or the `-copy-ownership-from` parameter with
the `-auto` parameter. When possible, use the `-auto` parameter rather
than the `-all` parameter to conform to disk ownership best practices. The
`-auto` parameter is ignored for array LUNs.

``[-p, -pool <integer>]`` - Pool::
This optional parameter specifies the pool to which a disk must be
assigned. It can take values of "0" or "1".

{ ``[-o, -owner <nodename>]`` - Owner Name::
This optional parameter specifies the node to which the disk or array LUN
has to be assigned.

``[-s, -sysid <nvramid>]`` - New Owner ID::
This optional parameter specifies the serial number (NVRAM ID) of the node
to which the disk or array LUN has to be assigned.

| ``[-copy-ownership-from <disk path name>]`` - Disk Name to Copy
Ownership }::
This optional parameter specifies the disk name from where the node needs
to copy disk ownership information. You can use this parameter for disks
to have the same ownership as the provided input disk.

``[-c, -checksum {block|zoned|advanced_zoned}]`` - Checksum
Compatibility::
```

```
 This optional parameter is used to set the checksum type for a disk or an
array LUN. The possible values are ``_block_`` , ``_zoned_`` , and
``_advanced_zoned_`` . This operation will fail if the specified disk is
incompatible with the specified checksum type. A newly created aggregate
with zoned checksum array LUNs is assigned advanced zoned checksum (AZCS)
checksum type. AZCS checksum type provides more functionality than the
"version 1" zoned checksum type which has been supported in previous ONTAP
releases. Zoned checksum spare array LUNs added to an existing zoned
checksum aggregate continue to be zoned checksum. Zoned checksum spare
array LUNs added to an AZCS checksum type aggregate use the AZCS checksum
scheme for managing checksums. For some disks (e.g. FCAL, SSD, SAS disks),
the checksum type cannot be modified. For more information on modifying
the checksum type, refer to the "Physical Storage Management Guide".

``[-f, -force <true>]`` - Force Flag::
 This optional parameter forces the assignment of ownership of an already
owned disk to a node. This parameter could also be used to assign an array
LUN with a redundancy error, for example, if the array LUN is available on
only one path. For a disk which is part of a live aggregate, even
specification of the `-force` parameter would not force the assignment,
since it would be catastrophic.

``[-N, -node <nodename>]`` - Node Name (For Auto Assign)::
 This optional parameter is used with either the `-auto` or the `-all`
parameter. If used with the `-auto` parameter, all disks which are visible
to the node specified in the `-node` parameter and which are eligible for
autoassignment would be assigned to it. If used with the `-all` parameter,
all unowned disks or array LUNs visible to the node would be assigned to
it.

{ ``[-root <true>]`` - Root Partition of Root-Data or Root-Data1-Data2
Partitioned Disk [.advaced-priv]#(privilege: advanced)#::
This optional parameter assigns the root partition of a root-data/root-
data1-data2 partitioned disk. You cannot use this parameter with disks
that are part of a storage pool. The default value is ``_false_`` .

| ``[-data <true>]`` - Data Partition of Root-Data Partitioned Disk
[.advaced-priv]#(privilege: advanced)#::
This optional parameter assigns the data partition of a root-data
partitioned disk. You cannot use this parameter with disks that are part
of a storage pool. The default value is ``_false_`` .

| ``[-data1 <true>]`` - Data1 Partition of Root-Data1-Data2 Partitioned
Disk [.advaced-priv]#(privilege: advanced)#::
This optional parameter assigns the data1 partition of a root-data1-data2
partitioned disk. You cannot use this parameter with disks that are part
```

```
of a storage pool. The default value is ``_false_`` .

| ``[-data2 <true>]`` - Data2 Partition of Root-Data1-Data2 Partitioned
Disk [.advaced-priv]#(privilege: advanced)# }::
This optional parameter assigns the data2 partition of a root-data1-data2
partitioned disk. You cannot use this parameter with disks that are part
of a storage pool. The default value is ``_false_`` .

== Examples

 The following example assigns ownership of an unowned disk named ``_1_``
.``_1_`` .``_16_`` to a node named ``_node1_`` :
```

cluster1::> storage disk assign -disk 1.1.16 -owner node1

```
The following example assigns all unowned disks or array LUNs visible to a
node named ``_node1_`` to itself:
```

cluster1::> storage disk assign -all -node node1

```
The following example autoassigns all unowned disks (eligible for
autoassignment) visible to a node named ``_node1_`` to itself:
```

cluster1::> storage disk assign -auto -node node1

```
The following two examples show the working of the `-force` parameter with
a spare disk that is already owned by another system:
```

cluster1::> storage disk assign -disk 1.1.16 -owner node1
Error: command failed: Failed to assign disks. Reason: Disk 1.1.16 is
already owned.

```

```

cluster1::> storage disk assign -disk 1.1.16 -owner node1 -force
Success.

```
The following example assigns ownership of the set of unowned disks on
``_<stack>_`` ``_1_`` , to a node named ``_node1_`` :
```

cluster1::> storage disk assign -disk 1.* -owner node1

```
The following example assigns ownership of unowned disk ``_1_`` .``_1_``
.``_16_`` by copying ownership from disk ``_1_`` .``_1_`` .``_18_`` :
```

cluster1::> storage disk assign -disk 1.1.16
-copy-ownership-from 1.1.18

```
 The following example assigns all unowned disks visible to a node named
``_node1_`` by copying ownership from disk ``_1_`` .``_1_`` .``_18_`` :
```

cluster1::> storage disk assign -all -node node1
-copy-ownership-from 1.1.18

```
The following example assigns the root partition of disk ``_1_`` .``_1_``
.``_16_`` to node1.
```

cluster1::> storage disk assign -disk 1.1.16 -owner node1 -root true
-force true

```
The following example assigns the data partition of root-data partitioned
disk ``_1_`` .``_1_`` .``_16_`` to node1.
```

cluster1::> storage disk assign -disk 1.1.16 -owner node1 -data true
-force true

```
The following example assigns the data1 partition of root-data1-data2
partitioned disk ``_1_`` .``_1_`` .``_24_`` to node1.
```

cluster1::> storage disk assign -disk 1.1.24 -owner node1 -data1 true
-force true

```
The following example assigns the data2 partition of root-data1-data2
partitioned disk ``_1_`` .``_1_`` .``_24_`` to node1.z33
```

cluster1::> storage disk assign -disk 1.1.24 -owner node1 -data2 true
-force true

```
[[ID41bbb6692a0cce0c2b7be440ee491e6e]]
= storage disk fail
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Fail the file system disk

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk fail` command can be used to manually force a file
system disk to fail. It is used to remove a file system disk that may be
logging excessive errors and requires replacement. To unfail a disk, use
the xref:{relative_path}storage-disk-unfail.html[storage disk unfail]
command.

== Parameters

``-disk <disk path name>`` - Disk Name::
This parameter specifies the disk to be failed.

``[-i, -immediate <true>]`` - Fail immediately::
This parameter optionally specifies whether the disk is to be failed
immediately. It is used to avoid Rapid RAID Recovery and remove the disk
from the RAID configuration immediately. Note that when a file system disk
has been removed in this manner, the RAID group to which the disk belongs
enters degraded mode (meaning a disk is missing from the RAID group). If a
suitable spare disk is available, the contents of the disk being removed
are reconstructed onto that spare disk.

== Examples

The following example fails a disk named 1.1.16 immediately:
```

cluster1::> storage disk fail -disk 1.1.16 -i true
WARNING: The system will not prefail the disk and its contents will not be
copied to a replacement disk before being failed out. Do you want to
fail out the disk immediately? {y|n}: y

```
== Related Links
```

```
* xref:{relative_path}storage-disk-unfail.html[storage disk unfail]




[[IDc903f3bfe93e31fa8259c99f9e34a258]]
= storage disk reassign
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Change the default owner of all disks from one node to
another

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage disk reassign` is deprecated and may be removed in a future
release of ONTAP. Disk reassignment is no longer required as part of a
controller replacement procedure. For further information, see the latest
controller or NVRAM FRU replacement flyer for your system. This command
changes the ownership of all disks on a node to the ownership of another
node. Use this command only when a node has a complete failure (for
instance, a motherboard failure) and is replaced by another node. If the
node's disks have already been taken over by its storage failover partner,
use the `-force` parameter.

== Parameters

``-s, -homeid <nvramid>`` - Current Home ID [.advaced-priv]#(privilege:
advanced)#::
This specifies the serial number of the failed node.

``-d, -newhomeid <nvramid>`` - New Home ID [.advaced-priv]#(privilege:
advanced)#::
This specifies the serial number of the node that is to take ownership of
the failed node's disks.

``[-f, -force <true>]`` - Force [.advaced-priv]#(privilege: advanced)#::
This optionally specifies whether to force the reassignment operation. The
default setting is `false` .
```

```
== Examples

In the following example, a node named node0 and having serial number
12345678 has failed. Its disks have not been taken over by its storage
failover partner. A replacement node with serial number 23456789 was
installed and connected to node0's disk shelves. To assign node0's disks
to the new node, start the new node and run the following command:
```

cluster::*> storage disk reassign -homeid 12345678 -newhomeid 23456789
node0’s disks 1.1.11, 1.1.12, 1.1.13, 1.1.14, 1.1.15, 1.1.16, 1.1.23 and 1.1.24
were reassigned to new owner with serial number 23456789.

```
In the following example, a similar failure has occurred, except that
node0's disks have been taken over by its storage failover partner, node1.
A new node with serial number 23456789 has been installed and configured.
To assign the disks that previously belonged to node0 to this new node,
run the following command:
```

cluster::*> storage disk reassign -homeid 12345678 -newhomeid 23456789 -force true
node0’s disks 1.1.11, 1.1.12, 1.1.13, 1.1.14, 1.1.15, 1.1.16, 1.1.23 and 1.1.24
were reassigned to new owner with serial number 23456789.

```
[[IDaf3bf20168a6093a73161d929a671c42]]
= storage disk refresh-ownership
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Refresh the disk ownership information on a node

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command updates the disk ownership information for all the disks
attached to a node to the latest view for all the nodes in the cluster.
During normal operations, disk ownership is kept up to date automatically.
In certain circumstances, however, disk ownership must be updated
manually. If this is required, EMS messages will indicate that this
command should be run. If the -node parameter is provided, the disk
ownership information is updated only on the node specified.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If this parameter is provided, the disk ownership information is updated
only on the node specified.

== Examples

The following example refreshes the disk ownership information for all the
nodes in the cluster:
```

cluster1::> storage disk refresh-ownership

```
[[ID6915206a7fd22d9d1d4c5dc982b68f3f]]
= storage disk remove-reservation
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Removes reservation from an array LUN marked as foreign.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage disk remove-reservation` command removes persistent
reservation from a specified foreign array LUN.

== Parameters

``-disk <disk path name>`` - Disk Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the disk from which persistent reservation is to be
removed.

== Examples

The following example removes the persistent reservation from a disk named
node1:switch01:port.126L1.
```

cluster1::> storage disk remove-reservation -disk node1:switch01:port.126L1

```
[[IDe01df6a7687844b9470e1a4b02215906]]
= storage disk remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a spare disk

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk remove` command removes the specified spare disk from
the RAID configuration, spinning the disk down when removal is complete.

This command does not remove disk ownership information from the disk.
Therefore, if you plan to reuse the disk in a different storage system,
you should use the xref:{relative_path}storage-disk-
removeowner.html[storage disk removeowner] command instead. See the
"Physical Storage Management Guide" for the complete procedure.

NOTE: For systems with multi-disk carriers, it is important to ensure that
none of the disks in the carrier are filesystem disks before attempting
removal. To convert a filesystem disk to a spare disk, see
xref:{relative_path}storage-disk-replace.html[storage disk replace] .

== Parameters

``-disk <disk path name>`` - Disk Name::
This parameter specifies the disk to be removed.

== Examples

The following example removes a spare disk named 1.1.16:
```

cluster1::> storage disk remove -disk 1.1.16

```
== Related Links

* xref:{relative_path}storage-disk-removeowner.html[storage disk
removeowner]
* xref:{relative_path}storage-disk-replace.html[storage disk replace]




[[IDee5bbe2ae68936b2b8f70c9e3545b6d9]]
= storage disk removeowner
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove disk ownership

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk removeowner` command removes ownership from a specified
disk. Then disk can then be reassigned to a new owner.

== Parameters

``-disk <disk path name>`` - Disk Name::
This specifies the disk whose ownership is to be removed.

== Examples

The following example removes the ownership from a disk named 1.1.27.
```

cluster1::> storage disk removeowner -disk 1.1.27

```
[[ID2f84b24fc32042eb9c638dc82f112727]]
= storage disk replace
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Initiate or stop replacing a file-system disk

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk replace` command starts or stops the replacement of a
file system disk with spare disk. When you start a replacement, Rapid RAID
Recovery begins copying data from the specified file system disk to a
spare disk. When the process is complete, the spare disk becomes the
active file system disk and the file system disk becomes a spare disk. If
you stop a replacement, the data copy is halted, and the file system disk
and spare disk retain their initial roles.

== Parameters

``-disk <disk path name>`` - Disk Name::
This specifies the file system disk that is to be replaced. Disk names
take one of the following forms:
+
  * Disks are named in the form ``_<stack-id>_`` .``_<shelf>_``
.``_<bay>_``
  * Disks on multi-disk carriers are named in the form ``_<stack-id>_``
.``_<shelf>_`` .``_<bay>_`` .``_<lun>_``
  * Virtual disks are named in the form <prefix>.<number>, where prefix is
the storage array's prefix and number is a unique ascending number.

+
Disk names take one of the following forms on clusters that are not yet
fully upgraded to ONTAP 8.3:
+
  * Disks that are not attached to a switch are named in the form
``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_`` . For disks with a
LUN, the form is ``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_``
L``_<LUN>_`` . For instance, disk number 16 on host adapter 1a on a node
named node0a is named node0a:1a.16. The same disk on LUN lun0 is named
node0a:1a.16Llun0.
  * Disks that are attached to a switch are named in the form ``_<node>_``
:``_<switch_name>_`` :``_<switch_port>_`` .``_<loop_ID>_`` . For disks
with a LUN, the form is ``_<node>_`` :``_<switch_name>_``
:``_<switch_port>_`` .``_<loop_ID>_`` L``_<LUN>_`` . For instance, disk
number 08 on port 11 of switch fc1 on a node named node0a is named
node0a:fc1:11.08. The same disk on LUN lun1 is named
node0a:fc1:11.08Llun1.
```

```
+
Before the cluster is upgraded to ONTAP 8.3, the same disk can have
multiple disk names, depending on how the disk is connected. For example,
a disk known to a node named alpha as alpha:1a.19 can be known to a node
named beta as beta:0b.37. All names are listed in the output of queries
and are equally valid. To determine a disk's unique identity, run a
detailed query and look for the disk's universal unique identifier (UUID)
or serial number.

``-action {start | stop}`` - Action::
This specifies whether to start or stop the replacement process.

``[-replacement <disk path name>]`` - Replacement::
This specifies the spare disk that is to replace the file system disk.

``[-allow-same-carrier <true>]`` - Allow Same RAID Group Within Carrier::
This parameter can be used to allow two disks housed in the same carrier
to be in the same RAID group when you replace a disk in an aggregate.
+
Having disks in the same carrier in the same RAID group is not desirable
because a carrier failure can cause a simultaneous outage for two disks in
the same RAID group. You can replace a disk in an aggregate with a disk
that causes this situation, but when an alternate disk becomes available,
ONTAP automatically initiates a series of disk copy operations to put the
disks into different RAID groups. For this reason, you should use this
parameter only when necessary. When possible, ensure that disks housed in
the same carrier are in different RAID groups.
+
This parameter affects only the disk replace operation. It is not a
persistent attribute of the aggregate.

``[-m, -allow-mixing <true>]`` - Allow Mixing of Disks of Different RPM or
Pool::
This optional parameter specifies whether the disk can be replaced with
another disk of different RPM or from different Pool. This parameter
affects only the current disk replacement operation.

== Examples

The following example begins replacing a file system disk named
``_1.0.16_`` with a spare disk named ``_1.1.14_`` .
```

cluster1::> storage disk replace -disk 1.0.16 -replacement 1.1.14 -action start

```
[[ID2775e3ccb9d1d10f512663b872de5583]]
= storage disk sanitize-spare
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Sanitize a spare disk (non-cryptographically)

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage disk sanitize-spare` command sanitizes the specified spare
disk (non-cryptographically).

== Parameters

``-disk <disk path name>`` - Disk Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the spare disk to be sanitized (non-
cryptographically).

== Examples

The following example sanitize a spare disk named 1.1.16 (non-
cryptographically):
```

cluster1::> storage disk sanitize-spare -disk 1.1.16

```
[[ID00b2130ca5036fc340977261e771658a]]
= storage disk set-foreign-lun
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Sets or Unsets an array LUN as foreign

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk set-foreign-lun` command sets or unsets a specified
array LUN as foreign. This command will enable/disable the feature of
importing the data from foreign LUN.

== Parameters

``-disk <disk path name>`` - Disk Name::
This parameter specifies the array LUN which is to be set or unset as
foreign.

``-is-foreign-lun <true>`` - Is Foreign LUN::
If the parameter value specified is ``_true_`` then array LUN is set as
foreign. If the parameter value specified is ``_false_`` then array LUN
foreignness is cleared.

== Examples

The following example shows how to set an array LUN as foreign:
```

cluster1::> storage disk set-foreign-lun -disk EMC-1.1 -is-foreign-lun true

```
The following example shows how to mark an array LUN as not foreign:
```

cluster1::> storage disk set-foreign-lun -disk EMC-1.1 -is-foreign-lun false

```
[[ID2239ae13f2a344bf929cfcd45b405613]]
= storage disk set-led
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Identify disks by turning on their LEDs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk set-led` command controls the LED of a specified disk.

You can turn an LED on or off, cause it to blink or stop blinking, or test
it.

This command is useful for locating a disk in its shelf.

== Parameters

``-action {on|off|blink|blinkoff|testall|resetall}`` - Action::
This parameter specifies the state to which the LED is to be set. Possible
values include the following:
+
  * on - The LED is lit steadily
  * off - The LED is not lit
  * blink - The LED blinks
  * blinkoff - The LED stops blinking and is not lit
  * testall - This tests the operation of every disk enclosure's hardware
and drivers per node. Do not use this value in normal operation.
  * resetall - This resets the LED of every disk on the node and lights up
the LED of disks with faults.



{ ``[-disk <disk path name>]`` - Disk Name::
This specifies the disk whose LED is to be set. Disk names take one of the
following forms:
+
  * Disks are named in the form ``_<stack-id>_`` .``_<shelf>_``
.``_<bay>_``
  * Disks on multi-disk carriers are named in the form ``_<stack-id>_``
.``_<shelf>_`` .``_<bay>_`` .``_<lun>_``
  * Virtual disks are named in the form <prefix>.<number>, where prefix is
the storage array's prefix and number is a unique ascending number.

+
```

```
Disk names take one of the following forms on clusters that are not yet
fully upgraded to ONTAP 8.3:
+
  * Disks that are not attached to a switch are named in the form
``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_`` . For disks with a
LUN, the form is ``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_``
L``_<LUN>_`` . For instance, disk number 16 on host adapter 1a on a node
named node0a is named node0a:1a.16. The same disk on LUN lun0 is named
node0a:1a.16Llun0.
  * Disks that are attached to a switch are named in the form ``_<node>_``
:``_<switch_name>_`` :``_<switch_port>_`` .``_<loop_ID>_`` . For disks
with a LUN, the form is ``_<node>_`` :``_<switch_name>_``
:``_<switch_port>_`` .``_<loop_ID>_`` L``_<LUN>_`` . For instance, disk
number 08 on port 11 of switch fc1 on a node named node0a is named
node0a:fc1:11.08. The same disk on LUN lun1 is named
node0a:fc1:11.08Llun1.



+
Before the cluster is upgraded to ONTAP 8.3, the same disk can have
multiple disk names, depending on how the disk is connected. For example,
a disk known to a node named alpha as alpha:1a.19 can be known to a node
named beta as beta:0b.37. All names are listed in the output of queries
and are equally valid. To determine a disk's unique identity, run a
detailed query and look for the disk's universal unique identifier (UUID)
or serial number.

| ``[-adapter <text>]`` - Adapter Name::
The name of the adapter to which the shelves of disks of interest are
attached to.

``[-node {<nodename>|local}]`` - Node Name }::
The node for which action is to be taken.

``[-duration <integer>]`` - Duration (minutes)::
This specifies the duration, in minutes, that the LED is to remain in the
specified state. Only actions "on" and "blink" are supported.

``[-iteration <integer>]`` - Test iterations::
This specifies the number of iterations to run the action for. Only action
"test-all" is supported.

== Examples

The following example causes the LEDs on all disks whose names match the
pattern 1.0.* to turn on for 5 minutes:
```

Cluster1::> storage disk set-led -disk 1.0.* -action on -duration 5

```
The following example causes the LEDs on disks 1.0.0 and 1.0.1 to turn on
for 2 minutes:
```

Cluster1::> storage disk set-led -disk 1.0.0,1.0.1 -action on -duration 2

```
The following example causes the LEDs on all disks attached to adapter 0b
on Node2 to turn on for 1 minute:
```

Cluster1::> storage disk set-led -node Node2 -adapter 0b -action on -duration 1

```
The following example resets the LEDs on all disks on the local node and
causes the LEDs of disks with faults to turn on:
```

Cluster1::> storage disk set-led -action resetall

```
The following example tests the LEDs on all disks owned by the local node
for 3 iterations:
```

Cluster1::> storage disk set-led -action testall -iteration 3

```
[[IDb8a6d931900bfe9c8e9b3d629af73781]]
= storage disk show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of disk drives and array LUNs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk show` command displays information about disks and array
LUNs. Where it appears in the remainder of this document "disk" may refer
to either a disk or an array LUN. By default, the command displays the
following information about all disks in column style output:
```

```
* Disk name
* Usable space on the disk, in human readable units
* Shelf number
* Bay number
* Container type (aggregate, broken, foreign, labelmaint, maintenance,
mediator, remote, shared, spare, unassigned, unknown, volume, or
unsupported)
* Position (copy, data, dparity, orphan, parity, pending, present, shared
or tparity)
* Container name
* Owning node name

To display detailed information about a single disk, use the `-disk`
parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all disks, in column style output.

| ``[-broken  ]``::
Displays the following RAID-related information about broken disks:
+
  * Original owning node name
  * Checksum compatibility
  * Disk name
  * Outage reason
  * Host bus adapter
  * Shelf number
  * Bay number
  * Primary port / Channel
  * Pool
  * Disk type
  * RPM (Revolutions per minute)
  * Usable size in human readable units
  * Physical size in human readable units
  * Current owner node



| ``[-errors  ]``::
Displays the following disk information about the disks which have errors.
+
  * Disk Name
  * Error Type
  * Error Description and corresponding corrective action
```

```
| ``[-foreign  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following foreign LUN import related information about
foreign disks:
+
  * Disk name
  * Array name
  * Capacity in sectors
  * Capacity in mb
  * Serial Number



| ``[-longop  ]``::
Displays the following information about long-running disk operations, in
column style output:
+
  * Disk name
  * Whether the disk is marked as prefailed
  * Whether the disk is being replaced
  * Whether the disk is zeroed
  * Copy destination
  * Percentage of copy complete
  * Percentage of zeroing complete
  * Percentage of reconstruction complete



| ``[-maintenance  ]``::
Displays the following RAID-related information about disks in the
maintenance center:
+
  * Original owning node name
  * Checksum compatibility
  * Disk name
  * Outage Reason
  * Host bus adapter
  * Shelf number
  * Bay number
  * Primary port / Channel
  * Pool
  * Disk type
  * RPM (Revolutions per minute)
  * Usable size in human readable units
  * Physical size in human readable units
```

```
  * Current owner node



| ``[-ownership  ]``::
Displays the following ownership-related information:
+
  * Disk name
  * Aggregate name
  * Home node name
  * Owning node name
  * Disaster recovery home node name
  * Home node system id
  * Owning node system id
  * Disaster recovery home node system id
  * Reservation node system id
  * SyncMirror pool



| ``[-partition-ownership  ]``::
Displays the following ownership-related information for partitioned
disks:
+
  * Disk name
  * Aggregate name
  * Owner of root partition on a partitioned disk
  * Owner system id of root partition on a partitioned disk
  * Owner of data or data1 partition on a root-data or a root-data1-data2
partitioned disk respectively
  * Owner system id of data or data1 partition on a root-data or a root-
data1-data2 partitioned disk respectively
  * Owner of data2 partition on a root-data1-data2 partitioned disk
  * Owner system id of data2 partition on a root-data1-data2 partitioned
disk
  * Owner of the disk which is partitioned
  * Owner system id of the disk which is partitioned



| ``[-physical  ]``::
Displays the following information about the disk's physical attributes,
in column style output:
+
  * Disk name
  * Disk type
```

```
  * Disk vendor
  * Disk model
  * Firmware revision level
  * RPM (Revolutions per minute)
  * BPS (Bytes per sector)



| ``[-physical-location  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information about disks:
+
  * Disk name
  * Container type
  * Primary path
  * Location
  * Home node name
  * Physical size in human readable units



| ``[-port  ]``::
Displays the following path-related information:
+
  * Disk name and disk port associated with disk primary path
  * Disk name and disk port associated with the disk secondary path, for a
multipath configuration
  * Type, shelf, and bay information for the disks



| ``[-primary-paths  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information about disks:
+
  * Disk Name
  * Shelf
  * Bay
  * Container Type
  * Primary Path



| ``[-raid  ]``::
Displays the following RAID-related information:
+
  * Disk name
  * Container type (aggregate, broken, labelmaint, maintenance, mediator,
remote, shared, spare, unassigned, unknown, or volume)
```

```
  * Outage reason
  * Position (copy, data, dparity, orphan, parity, pending, present,
shared or tparity)
  * RAID group name
  * Aggregate name



| ``[-raid-info-for-aggregate  ]``::
Displays the following RAID-related information about the disks used in an
aggregate:
+
  * Owning node name
  * Aggregate name
  * Plex name
  * RAID group name
  * Position (copy, data, dparity, orphan, parity, pending, present,
shared or tparity)
  * Disk name
  * Host bus adapter
  * Shelf number
  * Bay number
  * Primary port / Channel
  * Pool
  * Disk type
  * RPM (Revolutions per minute)
  * Usable size in human readable units
  * Physical size in human readable units


+
When this parameter is specified, RAID groups that use shared disks are
not included. Use xref:{relative_path}storage-aggregate-show-
status.html[storage aggregate show-status] to show information for all
RAID groups and aggregates.

| ``[-spare  ]``::
Displays the following RAID-related information about available spare
disks:
+
  * Original owning node name
  * Checksum compatibility
  * Disk name
  * Host bus adapter
  * Shelf number
  * Bay number
```

```
  * Primary port / Channel
  * Pool
  * Disk type
  * Disk class
  * RPM (Revolutions per minute)
  * Usable size in human readable units
  * Physical size in human readable units
  * Current owner node



| ``[-ssd-wear  ]``::
Displays the following wear life related information about solid state
disks:
+
  * Rated Life Used : An estimate of the percentage of device life that
has been used, based on the actual device usage and the manufacturer's
prediction of device life. A value greater than 99 indicates that the
estimated endurance has been used, but this does not necessarily indicate
a device failure. Omitted if value is unknown.
  * Spare Blocks Consumed Limit : Spare blocks consumed percentage limit
reported by the device. When the Spare Blocks Consumed percentage for the
device reaches this read-only value, ONTAP initiates a disk copy operation
to prepare to remove the device from service. Omitted if value is unknown.
  * Spare Blocks Consumed : Percentage of device spare blocks that have
been used. Each device has a number of spare blocks that will be used when
a data block can no longer be used to store data. This value reports what
percentage of the spares have already been consumed. Omitted if value is
unknown.



| ``[-virtual-machine-disk-info  ]``::
Displays information about ONTAP virtual disks, their mapped datastores
and their specific backing device attributes, such as: disk or LUN,
adapter and initiator details (if applicable).
+
  * Disk name.
  * Name of the node.
  * ONTAP-supplied serial number of the system disk.
  * Size of the system disk.
  * Name of the disk backing store. A backing store represents a storage
location for virtual machine files. It can be a VMFS volume, a directory
on network-attached storage, or a local file system path.
  * File name of the virtual disk used by the hypervisor. Each ONTAP disk
is mapped to a unique VM disk file.
```

```
  * Type of the disk backing store. It can be a VMFS volume, a directory
on network-attached storage, or a local file system path.
  * Size of the disk backing store.
  * Full path to the backing store for network-attached storage. This
field is valid only for NAS connections.
  * Backing adapter PCI device ID for the virtual disk, for example
"50:00.0".
  * Backing adapter device name, for example "vmhba32".
  * Backing adapter model type, for example "LSI1064E".
  * Backing adapter driver name of the initiator.
  * The iSCSI name of the disk backing target. This field is valid only
for iSCSI connections.
  * The iSCSI IP address of the disk backing target. This field is valid
only for iSCSI connections.
  * SCSI device name for the backing disk. It takes the form target-
id:lun-id, for example "2:1".
  * Hypervisor-assigned unique ID of the backing device (disk or LUN).
  * Backing disk partition number where the corresponding VM disk file
resides.
  * Size of the backing device (disk or LUN).
  * Backing device manufacturer, for example "FUJITSU" or "IBM".
  * Backing device model, for example "MBE2073RC" or "LUN".
  * Storage account associated with the VM Disk.
  * Container associated with the VM Disk.
  * Page blob associated with the VM Disk.
  * Error (if any) while retrieving virtual disk details.



| ``[-vmdisk-backing-info  ]``::
Displays information about the backing disks on certain ONTAP-v models:
+
  * Disk name
  * Backing disk vendor
  * Backing disk model
  * Backing disk serial number
  * Backing disk device id



| ``[-instance  ]`` }::
Displays detailed disk information. If no disk path name is specified,
this parameter displays the same detailed information for all disks as
does the -disk parameter. If a disk path name is specified, then this
parameter displays the same detailed information for the specified disks
as does the -disk parameter.
```

```
``[-disk <disk path name>]`` - Disk Name::
Displays detailed information about the specified disks. Disk names take
one of the following forms:
+
  * Disks are named in the form ``_<stack-id>_`` .``_<shelf>_``
.``_<bay>_``
  * Disks on multi-disk carriers are named in the form ``_<stack-id>_``
.``_<shelf>_`` .``_<bay>_`` .``_<lun>_``
  * Virtual disks are named in the form <prefix>.<number>, where prefix is
the storage array's prefix and number is a unique ascending number.

+
Disk names take one of the following forms on clusters that are not yet
fully upgraded to ONTAP 8.3:
+
  * Disks that are not attached to a switch are named in the form
``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_`` . For disks with a
LUN, the form is ``_<node>_`` :``_<host_adapter>_`` .``_<loop_ID>_``
L``_<LUN>_`` . For instance, disk number 16 on host adapter 1a on a node
named node0a is named node0a:1a.16. The same disk on LUN lun0 is named
node0a:1a.16Llun0.
  * Disks that are attached to a switch are named in the form ``_<node>_``
:``_<switch_name>_`` :``_<switch_port>_`` .``_<loop_ID>_`` . For disks
with a LUN, the form is ``_<node>_`` :``_<switch_name>_``
:``_<switch_port>_`` .``_<loop_ID>_`` L``_<LUN>_`` . For instance, disk
number 08 on port 11 of switch fc1 on a node named node0a is named
node0a:fc1:11.08. The same disk on LUN lun1 is named
node0a:fc1:11.08Llun1.



+
Before the cluster is upgraded to ONTAP 8.3, the same disk can have
multiple disk names, depending on how the disk is connected. For example,
a disk known to a node named alpha as alpha:1a.19 can be known to a node
named beta as beta:0b.37. All names are listed in the output of queries
and are equally valid. To determine a disk's unique identity, run a
detailed query and look for the disk's universal unique identifier (UUID)
or serial number.

``[-owner {<nodename>|local}]`` - Owner::
Selects information about disks that are owned by the specified node.

``[-owner-id <nvramid>]`` - Owner System ID::
Selects the disks that are owned by the node with the specified system ID.

``[-is-foreign {true|false}]`` - Foreign LUN [.advaced-priv]#(privilege:
```

```
advanced)#::
Selects information about array LUNs that have been declared to be foreign
LUNs.

``[-uid <text>]`` - Disk Unique ID::
Selects the disks whose unique id matches this parameter value. A disk
unique identifier has the form:
`20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000 `

``[-aggregate <aggregate name>]`` - Aggregate::
Selects information about disks that belong to the specified aggregate.

``[-array-name <array name>]`` - Array Name::
 Selects information about the LUNs presented by the specified storage
array.

``[-average-latency <integer>]`` - Average I/O Latency Across All Active
Paths::
 Selects information about disks that have the specified average latency.

``[-bay <integer>]`` - Bay::
Selects information about disks that are located in the carrier within the
specified shelf bay.

``[-bps <integer>]`` - Bytes Per Sector::
 Selects information about disks that have the specified number of bytes
per sector. Possible settings are 512, 520, 4096, and 4160.

``[-carrier-id <text>]`` - Carrier ID::
Selects information about disks that are located within the specified
multi-disk carrier.

``[-checksum-compatibility {advanced_zoned | block | none}]`` - Checksum
Compatibility::
Selects information about disks that have the specified checksum
compatibility.

``[-class {capacity | performance | archive | solid-state | array |
virtual}]`` - Disk Class::
Selects information about disks that have the specified disk class.
+
  * capacity = Capacity-oriented, near-line disk types. Includes disk
types FSAS, BSAS and ATA.
  * performance = Performance-oriented, enterprise class disk types.
Includes disk types FCAL and SAS.
```

```
  * archive = Archive class SATA disks in multi-disk carrier storage
shelves. Includes disk type MSATA.
  * solid-state = Solid-state drives. Includes disk type SSD, SSD-CAP and
SSD-NVM.
  * array = Logical storage devices backed by storage arrays and used by
ONTAP as disks. Includes disk type LUN.
  * virtual = Virtual disks that are formatted and managed by the
hypervisor. Includes disk type VMDISK.



``[-container-type {aggregate | broken | foreign | labelmaint |
maintenance | mediator | remote | shared | spare | unassigned | unknown |
unsupported}]`` - Container Type::
Selects information about disks that have the specified container type.
+
  * Aggregate = Disk is used as a physical disk in an aggregate.
  * Broken = Disk is in broken pool.
  * Foreign = Array LUN has been marked foreign.
  * Labelmaint = Disk is in online label maintenance list.
  * Maintenance = Disk is in maintenance center.
  * Mediator = A mediator disk is a disk used on non-shared HA systems
hosted by an external node which is used to communicate the viability of
the storage failover between non-shared HA nodes.
  * Remote = Disk belongs to the remote cluster.
  * Shared = Disk is partitioned or in a storage pool.
  * Spare = Disk is a spare disk.
  * Unassigned = Disk ownership has not been assigned.
  * Unknown = Container is currently unknown. This is the default setting.
  * Unsupported = Disk is not supported.



``[-container-name <text>]`` - Container Name::
Selects information about disks that have the specified container name. +
If a disk is in an aggregate or storage pool, the container name is the
name of the aggregate or storage pool. + Spare disks show the SyncMirror
Pool to which they belong. + Partitioned disks could return multiple
aggregate names.

``[-copy-destination <disk path name>]`` - Copy Destination Name::
Selects information about disks whose contents are being copied (due to
either Rapid RAID Recovery or disk replacement) to the specified spare
disk.

``[-copy-percent <integer>]`` - Percentage of Copy Complete::
```

```
Selects information about disks that are involved as either a source or
destination of a copy operation, (due to either disk replacement or Rapid
RAID Recovery) and that have the specified percentage of the copy
operation completed.

``[-data-owner {<nodename>|local}]`` - Owner of Data Partition of Root-
Data Partitioned Disk::
Selects information about disks that have the specified data partition
owner name. Used with root-data partitioned disks.

``[-data1-owner {<nodename>|local}]`` - Owner of Data1 Partition of Root-
Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified data1 partition
owner name. Used with root-data1-data2 partitioned disks.

``[-data2-owner {<nodename>|local}]`` - Owner of Data2 Partition of Root-
Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified data2 partition
owner name. Used with root-data1-data2 partitioned disks.

``[-data-home {<nodename>|local}]`` - Home Owner of Data Partition of
Root-Data Partitioned Disk::
Selects information about disks that have the specified data partition
home owner name. Used with root-data partitioned disks.

``[-data1-home {<nodename>|local}]`` - Home Owner of Data1 Partition of
Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified data1 partition
home owner name. Used with root-data1-data2 partitioned disks.

``[-data2-home {<nodename>|local}]`` - Home Owner of Data2 Partition of
Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified data2 partition
home owner name. Used with root-data1-data2 partitioned disks.

``[-data-owner-id <nvramid>]`` - Owner System ID of Data Partition of
Root-Data Partitioned Disk::
Selects information about disks that have the specified data partition
owner system ID. Used with root-data partitioned disks.

``[-data1-owner-id <nvramid>]`` - Owner System ID of Data1 Partition of
Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified data1 partition
owner system ID. Used with root-data1-data2 partitioned disks.

``[-data2-owner-id <nvramid>]`` - Owner System ID of Data2 Partition of
Root-Data1-Data2 Partitioned Disk::
```

```
Selects information about disks that have the specified data2 partition
owner system ID. Used with root-data1-data2 partitioned disks.

``[-data-home-id <nvramid>]`` - Home Owner System ID of Data Partition of
Root-Data Partitioned Disk::
Selects information about disks that have the specified data partition
home owner system ID. Used with root-data partitioned disks.

``[-data1-home-id <nvramid>]`` - Home Owner System ID of Data1 Partition
of Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified data1 partition
home owner system ID. Used with root-data1-data2 partitioned disks.

``[-data2-home-id <nvramid>]`` - Home Owner System ID of Data2 Partition
of Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified data2 partition
home owner system ID. Used with root-data1-data2 partitioned disks.

``[-disk-io-kbps-total <integer>]`` - Total Disk Throughput in KBPS Across
All Active Paths::
Selects information about disks that have attained the specified I/O
throughput on all connected paths.

``[-disk-iops-total <integer>]`` - Total Disk IOPs Across All Active
Paths::
Selects information about disks that have achieved the specified number of
IOPs per second on all connected paths.

``[-diskpathnames <disk path name>,...]`` - List of Path-Based Disk
Names::
Selects information about disks that have all of the specified path names.

``[-effective-rpm <integer>]`` - Effective RPM::
Selects information about disks with the specified effective rotational
speed.

``[-dr-home {<nodename>|local}]`` - Disaster Recovery Home::
Selects information about disks that have the specified Disaster home
node.

``[-dr-home-id <nvramid>]`` - Disaster Recovery Home System ID::
Selects information about disks whose Disaster home node has the specified
system id.

``[-drawer <integer>]`` - Drawer::
Selects information about disks that are located in the specified drawer.
```

```
``[-error-type
{onepath|onedomain|control|foreign|toobig|toosmall|invalidblocksize|target
asymmap|deviceassymmap|failovermisconfig|unknown|netapp|fwdownrev|qualfail
|diskfail|notallflashdisk}]`` - Error Type::
Selects information about disks that have the specified error types.
+
  * onepath = The array LUN is accessible only via a single path.
  * onedomain = The array LUN is accessible only via a single fault
domain.
  * control = The array LUN cannot be used because it is a control device.
  * foreign = The array LUN is marked as foreign and has some external
SCSI reservations other than those from ONTAP.
  * toobig = The array LUN exceeds the maximum array LUN size that ONTAP
supports.
  * toosmall = The array LUN is less than the minimum array LUN size that
ONTAP supports.
  * invalidblocksize = The array LUN is not a valid block size.
  * targetasymmap = The array LUN is presented more than once on a single
target port.
  * deviceassymmap = The array LUN is presented with multiple IDs.
  * failovermisconfig = The array LUN is configured with inconsistent
failover methods.
  * unknown = The array LUN from a storage array that is not supported by
this version of ONTAP.
  * netapp = A SAN front-end LUN from one ONTAP system that is presented
as external storage to another ONTAP system.
  * fwdownrev = The disk firmware is a down version.
  * qualfail = The disk is not supported.
  * diskfail = The disk is in a failed state.
  * notallflashdisk = The disk does not match the All-Flash Optimized
personality of the system.



``[-firmware-revision <text>]`` - Firmware Revision::
Selects information about disks that have the specified firmware revision
level.

``[-home {<nodename>|local}]`` - Home::
Selects information about disks that have the specified home node.

``[-home-id <nvramid>]`` - Home System ID::
Selects information about disks whose home node has the specified system
ID.

``[-host-adapter <text>]`` - Primary Path Host Adapter::
```

```
Selects information about disks that are currently using the specified
Host Bus Adapter.

``[-import-in-progress {true|false}]`` - Foreign LUN import in progress::
Selects information about the array LUNs that are currently being imported

``[-initiator <text>,...]`` - Initiator Port::
Selects information about disks that are visible to the initiator
specified. Disks that are not currently in use by that initiator are
included.

``[-initiator-iops <integer>,...]`` - Number of IOPS on Initiator (Rolling
Average)::
Selects information about disks that are visible to an initiator that has
executed the specified number of IOPs.

``[-initiator-io-kbps <integer>,...]`` - Kbytes of I/O per second on
Initiator (Rolling Average)::
Selects information about disks visible to an initiator that has executed
I/O at the specified throughput.

``[-initiator-lun-in-use-count <integer>,...]`` - Number of LUNs in the
in-use state on this initiator::
Selects information about disks with a path through an initiator that has
the specified in-use-count.

``[-initiator-side-switch-port <text>,...]`` - Initiator Side Switch
Port::
Selects information about disks that are visible to an initiator connected
to the specified switch port.

``[-is-multidisk-carrier {true|false}]`` - Multi Disk Carrier?::
Selects information about disks that are located within a multi-disk
carrier.

``[-is-local-attach {true|false}]`` - Indicates If the Disk Is Local to
This Cluster::
Selects information about disks attached to the local(true) or
remote(false) MetroCluster site.

``[-location {<nodename>|local}]`` - Physical Location::
Selects information about disks attached to the specified node.

``[-location-id <nvramid>]`` - The system ID of the node where the disk is
attached::
Selects information about disks attached to the node with the specified
system ID.
```

```
``[-lun <integer>,...]`` - LUN ID::
Selects information about the specified LUNs.

``[-lun-iops <integer>,...]`` - Number IOPS per second on disk (Rolling
Average)::
Selects information about the LUNs that have reached the specified number
of IOPs.

``[-lun-io-kbps <integer>,...]`` - Kbytes/sec on Disk (Rolling Average)::
Selects information about the LUNs that have reached the specified I/O
throughput.

``[-lun-path-use-state <text>,...]`` - The Use State of the LUN on this
path::
Selects information about LUNs reporting the specified in-use state.

``[-model <text>]`` - Model::
Selects information about disks of the specified model.

``[-nodelist {<nodename>|local}]`` - Controller name::
Selects information about disks that are visible to all of the specified
nodes .

``[-outage-reason <text>]`` - Outage Reason::
Selects information about disks that are not in service for the specified
reason. Possible values are: admin failed, admin removed, admin testing,
evacuated, bad label, bypassed, failed, init failed, label version,
labeled broken, labelmaint, LUN resized, missing, not responding, predict
failure, rawsize shrank, recovering, sanitizing, sanitized, SnapLock Disk,
testing, unassigned, unknown.

``[-path-error-count <integer>]`` - Path Error Count::
Selects information about disks that are visible on a path that has
incurred the specified number of errors.

``[-path-iops <integer>,...]`` - Number of IOPS on Path (Rolling
Average)::
Selects information about disks on those paths that have reached the
specified number of IOPs.

``[-path-io-kbps <integer>,...]`` - Kbytes of I/O per second on Path
(Rolling Average)::
Selects information about disk with paths that have reached the specified
I/O throughput

``[-path-link-errors <integer>,...]`` - Link Error count on path::
```

```
Selects information about disks with paths that have incurred the
specified number of FC link errors.

``[-path-lun-in-use-count <integer>,...]`` - Number of LUNs in the in-use
state on this path::
Selects information about disks with paths that have the specified in-use-
count.

``[-path-quality <integer>,...]`` - Percentage of weighted error
threshold::
Selects information about disks on paths that have incurred the specified
number of errors. The value displayed is a measure of the health of a path
expressed as a percentage of an error threshold. Once a path has reached
or surpassed the error threshold, another path will be selected for I/O
transfer, if there is one available.

``[-physical-size-mb <integer>]`` - Physical Size (MB)::
Selects information about disks that have the specified physical capacity,
in megabytes.

``[-physical-size {<integer>[KB|MB|GB|TB|PB]}]`` - Physical Size::
Selects information about disks that have the specified physical capacity,
in human readable units.

``[-physical-size-512b <integer>]`` - Physical Size in Units of 512
Bytes::
Selects information about disks that have the specified physical capacity,
in 512-byte chunks. This parameter is present only for backwards
compatibility with ONTAP 8.0.

``[-plex <text>]`` - Plex Name::
Selects information about disks that belong to the specified RAID plex.

``[-pool <text>]`` - Assigned Pool::
Selects information about disks that belong to the specified SyncMirror
pool (pool0 or pool1).

``[-port-speed <text>,...]`` - Port Speed::
Selects information about disks that are served by a Host Bus Adapter that
is running at the specified port speed.

``[-position <diskPositionType>]`` - Disk Position::
Selects information about disks that have the specified position within
their disk container.

``[-power-on-hours <integer>]`` - Hours Powered On::
Selects information about disks that have the specified number of hours
```

```
being powered up.

``[-prefailed {true|false}]`` - Marked for Rapid RAID Recovery?::
Selects information about disks that match the specified parameter value
indicating whether the disk is either awaiting or is in process of Rapid
RAID Recovery.

``[-preferred-target-port {true|false}]`` - Whether or not target port
group is preferred [.advaced-priv]#(privilege: advanced)#::
Selects information about disks that match the specified parameter value
indicating whether the backing storage is ALUA (Asymmetric Logical Unit
Access) capable and has specified the array target port on this path to be
a preferred target port for I/O.

``[-primary-port <text>]`` - Primary Path Disk Port::
Selects information about disks that use the specified primary port.

``[-raid-group <text>]`` - Raid Group Name::
Selects information about disks that belong to the specified RAID group.

``[-reconstruction-percent <integer>]`` - Percentage of Reconstruction
Complete::
Selects information about disks that are being reconstructed and that have
the specified percentage of the reconstruction operation completed.

``[-replacing {true|false}]`` - Being Replaced?::
Selects information about disks that match the specified boolean value
indicating whether the disk is either awaiting or in process of disk
replacement.

``[-reservation-key <text>]`` - Reservation Key::
If this parameter is specified, the command displays information only
about the disk or disks that have the specified persistent reservation
key.

``[-reservation-type {rs|we|re|ea|sa|wero|earo|wear|eaar|none}]`` -
Reservation Type::
If this parameter is specified, the command displays information only
about the disk or disks that have the specified persistent reservation
type. Possible values are: rs, we, re, ea, sa, wero, earo, wear, eaar, or
none.

``[-reserver-id <integer>]`` - Reservation System ID::
Selects information about disks that are reserved by the node with the
specified system ID.

``[-root-owner {<nodename>|local}]`` - Owner of Root Partition of Root-
```

```
Data/Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified root partition
owner name. Used with root-data/root-data1-data2 partitioned disks.

``[-root-owner-id <nvramid>]`` - Owner System ID of Root Partition of
Root-Data/Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified root partition
owner system ID. Used with root-data/root-data1-data2 partitioned disks.

``[-root-home {<nodename>|local}]`` - Home Owner of Root Partition of
Root-Data/Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified root partition
home owner name. Used with root-data/root-data1-data2 partitioned disks.

``[-root-home-id <nvramid>]`` - Home Owner System ID of Root Partition of
Root-Data/Root-Data1-Data2 Partitioned Disk::
Selects information about disks that have the specified root partition
home owner system ID. Used with root-data/root-data1-data2 partitioned
disks.

``[-rpm <integer>]`` - Revolutions Per Minute::
Selects information about disks that have the specified rotational speed.

``[-secondary-name <disk path name>]`` - Secondary Path Name::
Selects information about disks that use the specified secondary path
name, for multipath configuration.

``[-secondary-port <text>]`` - Secondary Path Disk Port::
Selects information about disks that use the specified secondary port.

``[-serial-number <text>]`` - Serial Number::
Selects information about the disk that has the specified serial number.

``[-storage-pool <text>]`` - Storage Pool Name::
Selects information about disks that belong to the specified SSD storage
pool.

``[-shelf <integer>]`` - Shelf::
Selects information about disks that are located within the specified
shelf.

``[-shelf-uid <text>]`` - Shelf UID::
Selects information about disks that are located within a shelf with the
specified Shelf UID.

``[-slot <integer>]`` - Slot::
Selects information about disks that are located in a drawer with the
```

```
specified slot.

``[-stack-id <integer>]`` - Stack ID::
A cluster unique id for a collection of one or more interconnected
shelves.

``[-target-iops <integer>,...]`` - Number of IOPS to Target (Rolling
Average)::
Selects information about disks that are visible on target ports that have
performed the specified number of IOPs.

``[-target-io-kbps <integer>,...]`` - Kbytes of I/O per second to Target
(Rolling Average)::
Selects information about disks that are visible on target ports that have
reached the specified I/O throughput.

``[-target-lun-in-use-count <integer>,...]`` - Number of LUNs in the in-
use state on this target::
Selects information about disks with a path through a target port that has
the specified in-use-count.

``[-target-port-access-state <text>,...]`` - Failover optimization type::
Selects information about disks that are visible on target ports that have
the specified access state.

``[-target-side-switch-port <text>,...]`` - Target Side Switch Port::
Selects information about disks that are visible on target ports
identified by the switch port to which they are connected.

``[-target-wwpn <text>,...]`` - Target Port::
Selects information about disks that are visible on target ports
identified by their World Wide Port Name.

``[-tpgn <integer>,...]`` - Target Port Group Number::
Selects information about disks that belong to the specified Target Port
Group Number.

``[-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS  | SSD | VMDISK |
SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Disk Type::
Selects information about disks that have the specified disk type.

``[-usable-size-mb <integer>]`` - Usable Size (MB)::
Selects information about disks that have the specified usable space, in
megabytes.

``[-usable-size {<integer>[KB|MB|GB|TB|PB]}]`` - Usable Size::
Selects information about disks that have the specified usable space, in
```

```
human readable units.

``[-vendor <text>]`` - Vendor Name::
Selects information about disks that have the specified vendor.

``[-vmdisk-device-id <integer>,...]`` - Virtual Disk Device ID::
Selects information about disks that have the specified virtual disk
device ID.

``[-zeroed {true|false}]`` - Zeroed?::
Selects information about disks that have (true) or have not (false) been
fully pre-zeroed.

``[-zeroing-percent <integer>]`` - Percentage of Zeroing Complete::
Selects information about disks that are zeroing and have the specified
percentage complete.

``[-carrier-serialno <text>]`` - Carrier Serial Number::
Selects information about disks that are located within the multi-disk
carrier specified by the serial number.

``[-vmdisk-target-address <text>]`` - Target Address of the VM Disk::
Displays the VM Disk's target address either in the form of bus target lun
or bus unit.

``[-path-count <integer>]`` - Number of Paths to Disk::
Selects information about disks that have the specified number of paths.

``[-hw-minimum-os <ClusterVersion>]`` - Hardware Minimum Supported ONTAP
Version (x.y.z)::
Selects information about disks that have the specified hardware minimum
supported ONTAP version in x.y.z format, for example "9.8.0" or "9.9.1".

``[-bridge-serialno <text>]`` - Bridge Serial Number::
Selects information about disks that are specified by the bridge serial
number.

== Examples

The following example displays information about all disks:
```

cluster1::> storage disk show
Usable Container
Disk Size Shelf Bay Type Position Aggregate Owner
--- --- --- --- --- --- --- ---
1.1.1 10GB 1 1 spare present - node1
1.1.4 78.59GB 1 4 spare present - node1

1.1.12 10GB 1 12 spare present - node1
1.2.12 10GB 2 12 broken present - node1
1.3.7 78.59GB 3 7 aggregate parity aggr0_u23 node1
1.1.6 78.59GB 1 6 broken present - node1
1.2.10 78.59GB 2 10 aggregate dparity aggr0_u23 node1
1.4.9 78.59GB 4 9 aggregate data aggr0_u23 node1
1.1.0 10GB 1 0 aggregate dparity aggr0_u22 node2
1.4.1 10GB 4 1 aggregate data dp_degraded node2
1.1.2 10GB 1 2 spare present - node2
1.1.3 20GB 1 3 spare present - node2
1.4.4 20GB 4 4 spare present - node2
1.4.6 10GB 4 6 aggregate data dp_sdc node2
1.1.5 268.0GB 1 5 maintenance present - node2
1.3.0 10GB 3 0 aggregate parity aggr0_u22 node2
1.4.11 10GB 4 11 spare present - node2
1.4.13 20GB 4 13 broken present - node2

```
The following example displays detailed information about a disk named
1.0.75
```

cluster1::> storage disk show -disk 1.0.75
Disk: 1.0.75
Container Type: spare
Owner/Home: node2 / node2
DR Home: -
Stack ID/Shelf/Bay: 1 / 0 / 75
LUN: 0
Array: N/A
Vendor: NETAPP
Model: X267_HKURO500SSX
Serial Number: ZAKAS0GH
UID:
1FF17846:0A419201:9325845A:3ABD5075:00000000:00000000:00000000:00000000:00000000:00000000
BPS: 512
Physical Size: 10.15GB
Position: present
Checksum Compatibility: block
Aggregate: -
Plex: -
Paths:
LUN Initiator Side Target Side Link
Controller Initiator ID Switch Port Switch Port Acc Use Target Port TPGN Speed I/O KB/s IOPS
--- --- --- --- --- --- --- --- --- --- --- --- --- ---
node1 0d 0 N/A N/A AO INU 220a000a3384e4d2 21 2 Gb/S 0 0
node1 0c 0 N/A N/A AO RDY 2209000a3384e4d2 62 2 Gb/S 0 0
node2 0d 0 N/A N/A AO INU 2209000a3384e4d2 62 2 Gb/S 3 0

Errors:
-

```
The following example displays RAID-related information about disks used
in an aggregate:
```

```
cluster1::> storage disk show -raid-info-for-aggregate
Owner Node: node1
Aggregate: aggr0_node1_0
Plex: plex0
RAID Group: rg0
Usable Physical
Position Disk HA Shelf Bay Chan Pool Type RPM Size Size
-------- --------------------- ------------ ---- ------ ----- ------ -------- --------
data 2.11.2 2d 11 2 B Pool0 SAS 15000 9.77GB 9.93GB
dparity 2.11.0 2d 11 0 B Pool0 SAS 15000 9.77GB 9.93GB
parity 2.11.1 2d 11 1 B Pool0 SAS 15000 9.77GB 9.93GB
Owner Node: node2
Aggregate: a1
Plex: plex0
RAID Group: rg0
Usable Physical
Position Disk HA Shelf Bay Chan Pool Type RPM Size Size
-------- --------------------- ------------ ---- ------ ----- ------ -------- --------
data 2.1.8 2a 1 8 B Pool0 BSAS 7200 9.77GB 9.91GB
dparity 2.1.6 2a 1 6 B Pool0 BSAS 7200 9.77GB 9.91GB
parity 2.1.7 2a 1 7 B Pool0 BSAS 7200 9.77GB 9.91GB
Owner Node: node2
Aggregate: a1
Plex: plex0
RAID Group: rg1
Usable Physical
Position Disk HA Shelf Bay Chan Pool Type RPM Size Size
-------- --------------------- ------------ ---- ------ ----- ------ -------- --------
data 2.1.11 2a 1 11 B Pool0 BSAS 7200 9.77GB 9.91GB
dparity 2.1.9 2a 1 9 B Pool0 BSAS 7200 9.77GB 9.91GB
parity 2.1.10 2a 1 10 B Pool0 BSAS 7200 9.77GB 9.91GB
Owner Node: node2
Aggregate: aggr0
Plex: plex0
RAID Group: rg0
Usable Physical
Position Disk HA Shelf Bay Chan Pool Type RPM Size Size
-------- --------------------- ------------ ---- ------ ----- ------ -------- --------
data 2.1.5 2a 1 5 B Pool0 BSAS 7200 9.71GB 10.03GB
dparity 2.1.2 2a 1 2 B Pool0 BSAS 7200 9.71GB 10.03GB
parity 2.1.4 2a 1 4 B Pool0 BSAS 7200 9.71GB 10.03GB
12 entries were displayed.
```

```
The following example displays RAID-related information about spares:
```

```
cluster1::> storage disk show -spare
Original Owner: node1
```

Checksum Compatibility: block
Usable Physical
Disk HA Shelf Bay Chan Pool Type RPM Size Size Owner
--- --- --- --- --- --- --- --- ---
1.1.23 0b 1 23 A Pool0 FCAL 10000 132.8GB 134.2GB node1
1.1.25 0b 1 25 A Pool0 FCAL 10000 132.8GB 133.9GB node1
1.1.26 0b 1 26 A Pool1 FCAL 10000 132.8GB 133.9GB node1
1.1.27 0b 1 27 A Pool1 FCAL 10000 132.8GB 134.2GB node1
Home Owner: node2
Checksum Compatibility: block
Usable Physical
Disk HA Shelf Bay Chan Pool Type RPM Size Size Owner
--- --- --- --- --- --- --- --- ---
1.1.19 0a 1 19 B Pool1 FCAL 10000 132.8GB 133.9GB node2
1.1.20 0a 1 20 B Pool0 FCAL 10000 132.8GB 133.9GB node2
1.1.21 0a 1 21 B Pool0 FCAL 10000 132.8GB 133.9GB node2

```
The following example displays RAID-related information about broken
disks:
```

cluster1::> storage disk show -broken
Original Owner: node1
Checksum Compatibility: block
Usable Physical
Disk Outage Reason HA Shelf Bay Chan Pool Type RPM Size Size
--- --- --- --- --- --- --- --- ---
1.1.0 admin failed 0b 1 0 A Pool0 FCAL 10000 132.8GB 133.9GB
1.2.6 admin removed 0b 2 6 A Pool1 FCAL 10000 132.8GB 134.2GB
Original Owner: node2
Checksum Compatibility: block
Usable Physical
Disk Outage Reason HA Shelf Bay Chan Pool Type RPM Size Size
--- --- --- --- --- --- --- --- ---
1.1.0 admin failed 0a 1 0 B Pool0 FCAL 10000 132.8GB 133.9GB
1.1.13 admin removed 0a 1 13 B Pool0 FCAL 10000 132.8GB 133.9GB
4 entries were displayed.

```
The following example displays RAID-related information about disks in
maintenance center:
```

cluster1::> storage disk show -maintenance
Original Owner: node1
Checksum Compatibility: block
Usable Physical
Disk Outage Reason HA Shelf Bay Chan Pool Type RPM Size Size
--- --- --- --- --- --- --- --- ---
1.1.8 admin testing 0b 1 8 A Pool0 FCAL 10000 132.8GB 133.9GB
1.2.11 admin testing 0b 2 11 A Pool1 FCAL 10000 132.8GB 134.2GB
Original Owner: node2
Checksum Compatibility: block

Usable Physical
Disk Outage Reason HA Shelf Bay Chan Pool Type RPM Size Size
--- --- --- --- --- --- --- --- ---
1.2.10 admin testing 0a 2 10 B Pool1 FCAL 10000 132.8GB 133.9GB
1.2.13 admin testing 0a 2 13 B Pool1 FCAL 10000 132.8GB 134.2GB
4 entries were displayed.

```
The following example displays partition-related information about disks:
```

cluster1::> storage disk show -partition-ownership
Disk Partition Home Owner Home ID Owner ID
--- --- --- --- --- ---
VMw-1.13 Container pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
VMw-1.14 Container pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
VMw-1.15 Container pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Root pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Data pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
VMw-1.16 Container pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Root pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Data1 pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Data2 pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
VMw-1.17 Container pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
VMw-1.18 Container pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Root - - - -
Data pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
VMw-1.19 Container pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Root pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786
Data1 - - - -
Data2 pvaruncluster-2-01 pvaruncluster-2-01 4087518786 4087518786

```
== Related Links

* xref:{relative_path}storage-aggregate-show-status.html[storage aggregate
show-status]




[[ID1f58380a5f95e3d7587460ffbdb54ef9]]
= storage disk unfail
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Unfail a broken disk

*Availability:* This command is available to _cluster_ administrators at
```

```
the _advanced_ privilege level.

== Description

The `storage disk unfail` command can be used to unfail a broken disk.

If the attempt to unfail the disk is unsuccessful, the disk remains in the
broken state.

The disk unfail command prompts for confirmation unless you specify the "-
quiet" parameter.

== Parameters

``-disk <disk path name>`` - Disk Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the disk to be unfailed.

``[-s, -spare <true>]`` - Make the Disk Spare [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies whether the unfailed disk will be made a spare
disk. The disk is forced to become a spare disk if this parameter is
specified.
+
If this parameter is not specified, the disk is brought back into its
parent aggregate. Setting this parameter might result in the aggregate
coming back online if it is not complete or online. The default value is
false.

``[-q, -quiet <true>]`` - Confirmations off [.advaced-priv]#(privilege:
advanced)#::
You can set this parameter to true to suppress the confirmation messages.
However, before proceeding with the command, you should be aware that the
confirmation message contains important information about the effect of
unfailing a disk. This command cannot be reversed once it is invoked. The
default value is false.

== Examples

The following example unfails a disk named 1.1.16:
```

cluster1::*> storage disk unfail -disk 1.1.16

```
Warning: Failed disk "1.1.16" may have aggregate labels and file system
data present. In that case, this command will attempt to bring this disk
back into the aggregate with which this disk had formerly been associated
```

```
and preserve file system data. Are you sure you want to continue with disk
unfail? {y|n}:




[[ID890231e443a718260e445c9de7aae9b8]]
= storage disk updatefirmware
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED) - Update disk firmware

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future release of
ONTAP. Use the "link:storage-disk-firmware-update.html[storage disk
firmware update] " command.



The `storage disk updatefirmware` command updates the firmware on one or
more disks.

You can download the latest firmware by using the
xref:{relative_path}storage-firmware-download.html[storage firmware
download] command.

You can specify a list of one or more disks whose firmware is to be
updated by using the `-disk`

parameter, or you can update the firmware on all local disks by omitting
the `-disk` parameter.

== Parameters

``[-disk <disk path name>,...]`` - Disk::
This specifies the disk or disks whose firmware is to be updated.
+
If you do not specify this option, all local disks' firmware is updated.
```

```
== Examples

The following example updates the firmware on all disks:
```

cluster1::> storage disk updatefirmware

```
== Related Links

* xref:{relative_path}storage-disk-firmware-update.html[storage disk
firmware update]
* xref:{relative_path}storage-firmware-download.html[storage firmware
download]




[[ID922b00f185c31d2cac4172fcbabb4a27]]
= storage disk zerospares
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Zero non-zeroed spare disks

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk zerospares` command zeroes all non-zeroed spare disks in
all nodes or a specified node in the cluster. A node must be online to
zero disks. This operation must be done before a disk can be reused in
another aggregate. This version of ONTAP uses fast zeroing, which converts
a spare disk from non-zeroed to zeroed without the long wait times
required when physically zeroing a disk.

== Parameters

``[-owner {<nodename>|local}]`` - Owner::
If this parameter is specified, only non-zeroed spares assigned to the
specified node will be zeroed. Otherwise, all non-zeroed spares in the
cluster will be zeroed.

== Examples

The following example zeroes all non-zeroed spares owned by a node named
node4, using fast zeroing:
```

cluster1::> storage disk zerospares -owner node4

```
[[IDba36412163e4b6c9d7735a768103e88d]]
= storage disk error show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display disk component and array LUN configuration errors.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk error show` command displays disk component and array
LUN configuration errors.

== 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.

``[-uid <text>]`` - UID::
Displays the error information of the disk whose unique ID matches the
value you specify. A disk unique identifier has the form:
`20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:0
0000000:00000000 `

``[-array-name <array name>]`` - Array Name::
Displays the errors of the storage array whose name you specified.

``[-node {<nodename>|local}]`` - Controller Name::
Displays the error information for the disks on the clustered node whose
name you specified.

``[-disk <disk path name>]`` - Disk::
Displays detailed error information about the disk you specified.

``[-serial-number <text>]`` - Serial Number::
Displays the error information for the disk whose serial number you
```

```
specified.

``[-error-id <integer>,...]`` - Error ID::
Displays the error information for the disks whose Error IDs match IDs you
specified.

``[-error-type
{onepath|onedomain|control|foreign|toobig|toosmall|invalidblocksize|target
asymmap|deviceassymmap|failovermisconfig|unknown|netapp|fwdownrev|qualfail
|diskfail|notallflashdisk}]`` - Error Type::
Displays all disk errors of the error types you specified, grouped by
type.
+
  * onepath = The array LUN is accessible only via a single path.
  * onedomain = The array LUN is accessible only via a single fault
domain.
  * control = The array LUN cannot be used because it is a control device.
  * foreign = The array LUN is marked as foreign and has some external
SCSI reservations other than those from ONTAP.
  * toobig = The array LUN exceeds the maximum array LUN size that ONTAP
supports.
  * toosmall = The array LUN is less than the minimum array LUN size that
ONTAP supports.
  * invalidblocksize = The array LUN is not a valid block size.
  * targetasymmap = The array LUN is presented more than once on a single
target port.
  * deviceassymmap = The array LUN is presented with multiple IDs.
  * failovermisconfig = The array LUN is configured with inconsistent
failover methods.
  * unknown = The array LUN from a storage array that is not supported by
this version of ONTAP.
  * netapp = A SAN front-end LUN from one ONTAP system that is presented
as external storage to another ONTAP system.
  * fwdownrev = The disk firmware is a down version.
  * qualfail = The disk is not supported.
  * diskfail = The disk is in a failed state.
  * notallflashdisk = The disk does not match the All-Flash Optimized
personality of the system.



== Examples

The following example displays configuration errors seen in the system:
```

cluster1::> storage disk error show

Disk Error Type Error Text
--- --- ---
1.02.0 qualfail This disk failed dynamic disk qualification. Update the Disk Qaulification Package.

```
[[IDbb30be7074b6066ffcf75789492b1ca7]]
= storage disk firmware revert
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Revert disk firmware

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage disk firmware revert` command reverts firmware on all disks
or a specified list of disks on a node.

You can specify a list of one or more disks whose firmware is to be
reverted by using the `-disk` parameter.

You can revert the firmware on all the disks owned by a node by using the
`-node` parameter.

This command can make the disks inaccessible for up to five minutes after
the start of its execution. Therefore, the network sessions that use the
concerned node must be terminated before running the `storage disk
firmware revert` command. This is particularly true for CIFS sessions that
might be terminated when this command is executed.

If you need to view the current firmware versions, use the
xref:{relative_path}storage-disk-show.html[storage disk show]`-
fields```_firmware-revision_`` command. The following example displays
partial output from the xref:{relative_path}storage-disk-show.html[storage
disk show]`-fields```_firmware-revision_`` command, where the firmware
version for the disks is NA02:
```

cluster1::> storage disk show -fields firmware-revision
disk firmware-revision
--- ---
1.0.0 NA02
1.0.1 NA02
1.0.2 NA02

1.0.3 NA02
1.0.4 NA02
1.0.5 NA02

```
The firmware files are stored in the /mroot/etc/disk_fw directory on the
node. The firmware file name is in the form of "product-ID.revision.LOD".
For example, if the firmware file is for Seagate disks with product ID
X225_ST336704FC and the firmware version is NA01, the file name is
X225_ST336704FC.NA01.LOD. If the node in this example contains disks with
firmware version NA02, the /mroot/etc/disk_fw/X225_ST336704FC.NA01.LOD
file is downloaded to every disk when you execute this command.

How to Revert the Firmware for an HA Pair in a Cluster

Use the following procedure to perform a revert on the disks in an HA
environment:

* Make sure that the nodes are not in takeover or giveback mode.
* Download the latest firmware on both nodes by using the
xref:{relative_path}storage-firmware-download.html[storage firmware
download] command.
* Revert the disk firmware on Node A's disks by entering the `storage disk
firmware revert``-node```_node-A_`` command.
* Wait until the `storage disk firmware revert` command completes on Node
A, and then revert the firmware on Node B's disks by entering the `storage
disk firmware revert``-node```_node-B_`` command.



== Parameters

{ ``-disk <disk path name>,...`` - Disk Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the disk or disks whose firmware is to be reverted.

| ``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)# }::
Specifies the node name. The disk firmware will be reverted on all the
disks owned by the node specified by this parameter.

== Examples



* The following example reverts the firmware on all disks owned by
cluster-node-01:
```

cluster1::*> storage disk firmware revert -node cluster-node-01

Warning: Disk firmware reverts can be disruptive to the system. Reverts involve
power cycling all of the affected disks, as well as suspending disk
I/O to the disks being reverted. This delay can cause client
disruption. Takeover/giveback operations on a high-availability (HA)
group will be delayed until the firmware revert process is complete.
Disk firmware reverts should only be done one node at a time. Disk
firmware reverts can only be performed when the HA group is healthy;
they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware reverts? {y|n}: y

Info: Reverting disk firmware for disks on cluster-node-01.

```
* The following example reverts the firmware on disk 1.5.0 which is owned
by node cluster-node-04:
```

cluster1::*> storage disk firmware revert -disk 1.5.0

Warning: Disk firmware reverts can be disruptive to the system. Reverts involve
power cycling all of the affected disks, as well as suspending disk
I/O to the disks being reverted. This delay can cause client
disruption. Takeover/giveback operations on a high-availability (HA)
group will be delayed until the firmware revert process is complete.
Disk firmware reverts should only be done one node at a time. Disk
firmware reverts can only be performed when the HA group is healthy;
they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware reverts? {y|n}: y

Info: Reverting disk firmware for disks on cluster-node-04.

```
== Related Links

* xref:{relative_path}storage-disk-show.html[storage disk show]
* xref:{relative_path}storage-firmware-download.html[storage firmware
download]




[[ID6c234a5850c23dc376043e1679903bca]]
= storage disk firmware show-update-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display disk firmware update status.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage disk firmware show-update-status` command displays the state
of the background disk firmware update process.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - Node [.advaced-priv]#(privilege: advanced)#::
Selects the node that matches this parameter value.

``[-num-waiting-download <integer>]`` - The Number of Disks Waiting to
Download [.advaced-priv]#(privilege: advanced)#::
Selects the nodes whose number of disks waiting to download by the BDFU
process matches this parameter value.

``[-total-completion-estimate <integer>]`` - Estimated Duration to
Completion (mins) [.advaced-priv]#(privilege: advanced)#::
Selects the nodes whose Background Disk Firmware Update (BDFU) completion
time estimate matches this parameter value. This indicates the amount of
estimated time required for BDFU to complete the firmware update cycle.

``[-average-duration-per-disk <integer>]`` - Average Firmware Update
Duration per Disk (secs) [.advaced-priv]#(privilege: advanced)#::
Selects the nodes whose BDFU reports the average time required to update a
single disk matches this parameter value. This indicates the average
amount of time required by each disk drive.

``[-unable-to-update <disk path name>,...]`` - List of Disks with a Failed
Update [.advaced-priv]#(privilege: advanced)#::
Selects the nodes whose unable to update disk list matches this parameter
value. This is a list of disks that failed to update the firmware.
```

```
``[-update-status {off|running|idle}]`` - Background Disk Firmware Update
Status [.advaced-priv]#(privilege: advanced)#::
Selects the nodes whose BDFU process status matches this parameter value.
Possible values are:
+
  * off - The BDFU process is off.
  * running - The BDFU process is on and currently running.
  * idle - The BDFU process is on and is currently idle.



== Examples
```

cluster1::*> storage disk firmware show-update-status
Number Average Total
Update Waiting Duration Completion
Node State Download /Disk (Sec) Est. (Min) Unable to Update
---------- ------- -------- ---------- ---------- ----------
node1 running 2 120 4 1.3.3
node2 idle 0 120 0 -
node3 off 0 120 0 -

```
[[ID40cc9b286470f39f6b76ec836b16dc1d]]
= storage disk firmware update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Update disk firmware

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Use the `storage disk firmware update` command to manually update firmware
on all disks or a specified list of disks on a node. However, the
recommended way to update disk firmware in a cluster is to enable
automatic background firmware update by enabling the `-bkg-firmware-
update` parameter for all of the nodes in the cluster. You can do this by
entering the xref:{relative_path}storage-disk-option-modify.html[storage
disk option modify]`-node```_*_```-bkg-firmware-update```_on_`` command.

You can download the latest firmware on the node by using the
xref:{relative_path}storage-firmware-download.html[storage firmware
```

```
download] command.

You can specify a list of one or more disks whose firmware is to be
updated by using the `-disk` parameter.

You can update the firmware on all the disks owned by a node by using the
`-node` parameter.

This command can make the disks inaccessible for up to five minutes after
the start of its execution. Therefore, the network sessions that use the
concerned node must be terminated before running the `storage disk
firmware update` command. This is particularly true for CIFS sessions that
might be terminated when this command is executed.

The firmware is automatically downloaded to disks, which report previous
versions of the firmware. For information on automatic firmware update
downloads, see "Automatic versus Manual Firmware Download".

If you need to view the current firmware versions, use the
xref:{relative_path}storage-disk-show.html[storage disk show]`-
fields```_firmware-revision_`` command. The following example displays
partial output from the xref:{relative_path}storage-disk-show.html[storage
disk show]`-fields```_firmware-revision_`` command, where the firmware
version for the disks is NA01:
```

```
cluster1::> storage disk show -fields firmware-revision
disk firmware-revision
-------- ----------------
1.0.0 NA01
1.0.1 NA01
1.0.2 NA01
1.0.3 NA01
1.0.4 NA01
1.0.5 NA01
```

```
The firmware files are stored in the /mroot/etc/disk_fw directory on the
node. The firmware file name is in the form of "product-ID.revision.LOD".
For example, if the firmware file is for Seagate disks with product ID
X225_ST336704FC and the firmware version is NA02, the filename is
X225_ST336704FC.NA02.LOD. The revision part of the file name is the number
against which the node compares each disk's current firmware version. If
the node in this example contains disks with firmware version NA01, the
/mroot/etc/disk_fw/X225_ST336704FC.NA02.LOD file is used to update every
eligible disk when you execute this command.

Automatic versus Manual Firmware Download
```

```
The firmware is automatically downloaded to those disks that report
previous versions of firmware following a system boot or disk insertion.
Note that:

* A manual download is a disruptive operation that makes disks
inaccessible for up to five minutes after the download is started. Network
sessions that use the node must be terminated before running the `storage
disk firmware update` command.
* The firmware is not automatically downloaded to the node's partner node
in an HA pair.
* The firmware is not automatically downloaded to unowned disks on nodes
configured to use software-based disk ownership.
* The bkg-firmware-update parameter controls how the automatic firmware
download feature works:

  * If the bkg-firmware-update parameter is set to off, then the `storage
disk firmware update` will update the firmware on the drives in parallel.
  * If the bkg-firmware-update parameter is set to on, then the `storage
disk firmware update` will update spares and filesystem disks in a
nondisruptive manner in the background after boot. Firmware downloads for
these disks will be done sequentially by temporarily taking them offline
one at a time for the duration of the download. After the firmware is
updated, the disk will be brought back online and restored to its normal
operation.



During an automatic download to an HA environment, the firmware is not
downloaded to the disks owned by the HA partner.

When you use the `storage disk firmware update` command, the firmware is:

* Updated on every disk regardless of whether it is on the A-loop, the B-
loop, or in an HA environment.
* If the node is configured in a software-based disk ownership system,
only disks owned by this node are updated.



During an automatic firmware download in a MetroCluster(TM) environment,
the firmware is not downloaded to the disks owned by the partner cluster.
During both manual and automatic firmware download in a MetroCluster-over-
IP environment, the firmware is not downloaded to any remote disks located
at the partner cluster while Disaster Recovery is in progress.

Follow the instructions in "How to Update the Firmware for an HA Pair in a
Cluster" to ensure that the updating process is successful. ONTAP supports
```

```
redundant path configurations for disks in a non-HA configuration. The
firmware is automatically downloaded to disks on the A-loop or B-loop of
redundant configurations that are not configured in an HA pair and are not
configured to use software-based disk ownership.

Automatic Background Firmware Update

The firmware can be updated in the background so that the firmware update
process does not impact the clients. This functionality is controlled with
the bkg-firmware-update parameter. You can modify the parameter by using
the CLI xref:{relative_path}storage-disk-option-modify.html[storage disk
option modify]`-node```_node_name_```-bkg-firmware-update```_on|off_``
command. The default value for this parameter is "on".

When disabled or set to "off", `storage disk firmware update` will update
the firmware in automated mode. This means that all disks which had older
firmware revision will be updated regardless of whether they are spare or
filesystem disks.

When enabled or set to "on", the background `storage disk firmware update`
will update firmware in automated mode only on disks that can be
successfully taken offline from active filesystem RAID groups and from the
spare pool. To ensure a faster boot process, the firmware is not
downloaded to spares and filesystem disks at boot time.

This provides the highest degree of safety available, without the cost of
copying data from each disk in the system twice. Disks are taken offline
one at a time and then the firmware is updated on them. The disk is
brought online after the firmware update and a mini/optimized reconstruct
happens for any writes, which occurred while the disk was offline.
Background disk firmware update will not occur for a disk if its
containing RAID group or the volume is not in a normal state (for example,
if the volume/plex is offline or the RAID group is degraded). However, due
to the continuous polling nature of background disk firmware update,
firmware updates will resume after the RAID group/plex/volume is restored
to a normal mode. Similarly, background disk firmware updates are
suspended for the duration of any reconstruction within the system.

How to Update the Firmware for an HA Pair in a Cluster

The best way to update the firmware in a cluster with HA pairs is to use
automatic background firmware update by enabling the option bkg-firmware-
update parameter for each node. Enable the `-bkg-firmware-update`
parameter on all the nodes by entering the xref:{relative_path}storage-
disk-option-modify.html[storage disk option modify]`-
node```_node_name_```-bkg-firmware-update```_on_`` command. Alternatively,
```

```
use the following procedure to successfully perform a manual update on the
disks in an HA environment:

* Make sure that the nodes are not in takeover or giveback mode.
* Download the latest firmware on both the nodes by using the
xref:{relative_path}storage-firmware-download.html[storage firmware
download] command.
* Install the new disk firmware on Node A's disks by entering the `storage
disk firmware update``-node```_node-A_`` command.
* Wait until the `storage disk firmware update` command completes on Node
A, and then install the new disk firmware on Node B's disks by entering
the `storage disk firmware update``-node```_node-B_`` command.



== Parameters

{ ``-disk <disk path name>,...`` - Disk [.advaced-priv]#(privilege:
advanced)#::
Specifies the disk or disks whose firmware is to be updated.

| ``-node {<nodename>|local}`` - node [.advaced-priv]#(privilege:
advanced)# }::
Specifies the node name. The disk firmware will be updated on all the
disks owned by the node specified by this parameter.

== Examples



* The following example updates the firmware on all disks owned by
cluster-node-01:
```

cluster1::*> storage disk firmware update -node cluster-node-01

Warning: Disk firmware updates can be disruptive to the system. Updates involve
power cycling all of the affected disks, as well as suspending disk
I/O to the disks being updated. This delay can cause client
disruption. Takeover/giveback operations on a high-availability (HA)
group will be delayed until the firmware update process is complete.
Disk firmware updates should only be done one node at a time. Disk
firmware updates can only be performed when the HA group is healthy;
they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware updates? {y|n}: y

Info: Updating disk firmware for disks on cluster-node-01.

```
* The following example updates the firmware on disk 1.5.0 which is owned
by node cluster-node-04:
```

cluster1::*> storage disk firmware update -disk 1.5.0

Warning: Disk firmware updates can be disruptive to the system. Updates involve
power cycling all of the affected disks, as well as suspending disk
I/O to the disks being updated. This delay can cause client
disruption. Takeover/giveback operations on a high-availability (HA)
group will be delayed until the firmware update process is complete.
Disk firmware updates should only be done one node at a time. Disk
firmware updates can only be performed when the HA group is healthy;
they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware updates? {y|n}: y

Info: Updating disk firmware for disks on cluster-node-04.

```
== Related Links

* xref:{relative_path}storage-disk-option-modify.html[storage disk option
modify]
* xref:{relative_path}storage-firmware-download.html[storage firmware
download]
* xref:{relative_path}storage-disk-show.html[storage disk show]




[[ID18c8273185f8e488df92e825f21385a9]]
= storage disk option modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify disk options

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk option modify` command modifies the background firmware
update setting, automatic copy setting, controls automatic disk assignment
of all disks assigned to a specified node, or modifies the policy of
```

```
automatic disk assignment of unowned disks.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node that owns the disks whose options are to
be modified.

``[-bkg-firmware-update {on|off}]`` - Background Firmware Update::
This parameter specifies whether firmware updates run as a background
process. The default setting is `on` , which specifies that firmware
updates to spare disks and file system disks is performed nondisruptively
via a background process. If the option is turned off, automatic firmware
updates occur at system startup or during disk insertion.

``[-autocopy {on|off}]`` - Auto Copy::
This parameter specifies whether data is to be automatically copied from a
failing disk to a spare disk in the event of a predictive failure. The
default setting is `on` . It is sometimes possible to predict a disk
failure based on a pattern of recovered errors that have occurred. In such
cases, the disk reports a predictive failure. If this option is set to on,
the system initiates Rapid RAID Recovery to copy data from the failing
disk to an available spare disk. When data is copied, the disk is marked
as failed and placed in the pool of broken disks. If a spare is not
available, the node continues to use the disk until it fails. If the
option is set to off, the disk is immediately marked as failed and placed
in the pool of broken disks. A spare is selected and data from the missing
disk is reconstructed from other disks in the RAID group. The disk does
not fail if the RAID group is already degraded or is being reconstructed.
This ensures that a disk failure does not lead to the failure of the
entire RAID group.

``[-autoassign {on|off}]`` - Auto Assign::
This parameter specifies whether automatic assignment of unowned disks is
enabled or disabled. The default setting is `on` . This parameter is used
to set both a node-specific and a cluster-wide disk option.

``[-autoassign-policy {default|bay|shelf|stack}]`` - Auto Assignment
Policy::
 This parameter defines the granularity at which auto assign should work.
This option is ignored if the -autoassign option is off. Auto assignment
can be done at the stack/loop, shelf, or bay level. The possible values
for the option are default, stack, shelf, and bay. The default value is
platform dependent. It is stack for all non-entry platforms and single-
node systems, whereas it is bay for entry-level platforms.
```

```
== Examples

The following example sets the background firmware update setting to on
for all disks belonging to a node named node0:
```

cluster1::> storage disk option modify -node node0 -bkg-firmware-update on

```
The following example shows how to enable auto assignment for the disks on
node1:
```

cluster1::> storage disk option modify -node node1 -autoassign on
cluster1::> storage disk option show
Node BKg. FW. Upd. Auto Copy Auto Assign Auto Assign Policy
--- --- --- --- ---
node1 on on on default
node2 on on off default
2 entries were displayed.

```
The following example shows how to modify the auto assignment policy on
node1:
```

cluster1::> storage disk option modify -node node1 -autoassign-policy bay
cluster1::> storage disk option show
Node BKg. FW. Upd. Auto Copy Auto Assign Auto Assign Policy
--- --- --- --- ---
node1 on on on bay
node2 on on off default
2 entries were displayed.

```
[[ID6eab34b091df546018f789e62a3aa2be]]
= storage disk option show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of disk options

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage disk option show` command displays the settings of the
```

```
following disk options:

* Background firmware update
* Automatic copying of data to a spare disk in the event of a predictive
failure
* Automatic assignment of disks
* Policy that governs automatic assignment of unowned disks



== 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>|local}]`` - Node::
Selects the node that owns the disks. If this parameter is not specified,
the command displays information about the disk options on all the nodes.

``[-bkg-firmware-update {on|off}]`` - Background Firmware Update::
Selects the disks that match this parameter value.

``[-autocopy {on|off}]`` - Auto Copy::
Selects the disks that match this parameter value.

``[-autoassign {on|off}]`` - Auto Assign::
Displays the auto assignment status of unowned disks. The default value is
`on` .

``[-autoassign-policy {default|bay|shelf|stack}]`` - Auto Assignment
Policy::
Selects the disks that match the automatic assignment policy value:
+
  * Default
  * Stack/loop
  * Shelf
  * Bay



== Examples
```

```
The following example displays disk-option settings for disks owned by all
nodes in the cluster:
```

cluster1::> storage disk option show
Node BKg. FW. Upd. Auto Copy Auto Assign Auto Assign Policy
--- --- --- --- ---
node0 on on on default
node1 on on on stack
node2 on on on bay
node3 on on on bay
4 entries were displayed.

```
:leveloffset: -1


= storage dqp commands

:leveloffset: +1


[[IDa5df37131d5a2bf7b3b6979e7cd8b526]]
= storage dqp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Disk Qualification Package details

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage dqp show` command displays information about the Disk
Qualification Package (DQP). The command displays the following
information:

* Node Name
* Package Date
* File Version
* File Name
* Drive Record Count
* Drive Alias Record Count
* Device Class Record Count
```

```
* System Class Record Count



== 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>|local}]`` - Node::
Displays DQP information on the specified node.

``[-package-date <text>]`` - Package Date::
Selects DQP information with the specified package date.

``[-file-version <text>]`` - File Version::
Selects DQP information with the specified file version.

``[-filename <text>]`` - File Name::
Selects DQP information with the specified file name.

``[-drive-records <integer>]`` - Drive Record Count::
Selects DQP information with the specified drive record count.

``[-alias-records <integer>]`` - Alias Record Count::
Selects DQP information with the specified drive alias record count.

``[-device-records <integer>]`` - Device Record Count::
Selects DQP information with the specified device class record count.

``[-system-records <integer>]`` - System Record Count::
Selects DQP information with the specified system class record count.

== Examples

The following example displays the brief version of Disk Qualification
Package details:
```

cluster1::> storage dqp show
Node Package Date Version

--- --- ---
cluster1-01 20180907 3.17
cluster1-02 20180907 3.17
2 entries were displayed.

```
The following example displays the -instance version of Disk Qualification
Package details:
```

cluster1::> storage dqp show -instance
Node: cluster1-01
Package Date: 20180907
Version: 3.17
File Name: qual_devices_v3
Drive Record Count: 626
Alias Record Count: 196
Device Record Count: 28
System Record Count: 3
Node: cluster1-02
Package Date: 20180907
Version: 3.17
File Name: qual_devices_v3
Drive Record Count: 626
Alias Record Count: 196
Device Record Count: 28
System Record Count: 3
2 entries were displayed.

```
:leveloffset: -1


= storage encryption commands

:leveloffset: +1


[[ID8c7963ab0931dbd009b59a2881b47dd3]]
= storage encryption disk destroy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cryptographically destroy a self-encrypting disk

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `storage encryption disk destroy` command cryptographically destroys a
self-encrypting disk (SED), making it incapable of performing I/O
operations. This command performs the following operations:

* Employs the inherent erase capability of SEDs to cryptographically
sanitize the disk
* Permanently locks the disk to prevent further data access
* Changes the data and FIPS authentication keys to random values that are
not recorded except within the SED.

Use this command with extreme care. The only mechanism to restore the disk
to usability (albeit without the data) is the xref:{relative_path}storage-
encryption-disk-revert-to-original-state.html[storage encryption disk
revert-to-original-state] operation that is available only on disks that
have the physical secure ID (PSID) printed on the disk label.

The destroy command requires you to enter a confirmation phrase before
proceeding with the operation.

The command releases the cluster shell after launching the operation.
Monitor the output of the xref:{relative_path}storage-encryption-disk-
show-status.html[storage encryption disk show-status] command for command
completion.

Upon command completion, remove the destroyed SED from the system.

== Parameters

``-disk <disk path name>`` - Disk Name::
This parameter specifies the name of the disk you want to
cryptographically destroy. See the man page for the `storage disk modify`
command for information about disk-naming conventions.

``[-force-all-states <true>]`` - Destroy All Matching Disks::
When this parameter is ``_false_`` or not specified, the operation
defaults to spare and broken disks only, as reported in the output of the
xref:{relative_path}storage-disk-show.html[storage disk show] command.
When you specify this parameter as ``_true_`` , it allows you to
cryptographically destroy all matching disk names regardless of their
state, including those in active use in aggregates. This allows a quick
destroy of all system disks if you use the `-disk` parameter with the
asterisk wildcard (*). If you destroy active disks, the nodes might not be
able to continue operation, and might halt or panic.

== Examples
```

```
The following command cryptographically destroys the disk 1.10.20:
```

cluster1::> storage encryption disk destroy 1.10.20

Warning: This operation will cryptographically destroy 1 spare or broken
self-encrypting disks on 1 node.
You cannot reuse destroyed disks unless you revert
them to their original state using the PSID value.
To continue, enter
destroy disk
:destroy disk

Info: Starting destroy on 1 disk.
View the status of the operation by using the
"storage encryption disk show-status" command.

cluster1::>

```
If you do not enter the correct confirmation phrase, the operation is
aborted:
```

cluster1::> storage encryption disk destroy 1.10.2*

Warning: This operation will cryptographically destroy 5 spare or broken
self-encrypting disks on 1 node.
You cannot reuse destroyed disks unless you revert
them to their original state using the PSID value.
To continue, enter
destroy disk
:yes
No disks destroyed.

cluster1::>

```
The following command quickly cryptographically destroys all system disks,
including those in active use in aggregates and shared devices:
```

cluster1::> storage encryption disk destroy -force-all-states -disk *

Warning: This operation will cryptographically destroy 96
self-encrypting disks on 4 nodes.
To continue, enter
destroy disk
:destroy disk

Info: Starting destroy on 96 disks.
View the status of the operation by using the
storage encryption disk show-status command.

cluster1::>

```
== Related Links

* xref:{relative_path}storage-encryption-disk-revert-to-original-
state.html[storage encryption disk revert-to-original-state]
* xref:{relative_path}storage-encryption-disk-show-status.html[storage
encryption disk show-status]
* xref:{relative_path}storage-disk-show.html[storage disk show]




[[ID171fc9e44b2e281104eaf9c3fcda9b3a]]
= storage encryption disk modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify self-encrypting disk parameters

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage encryption disk modify` command changes the data protection
parameters of self-encrypting disks (SEDs) and FIPS-certified SEDS (FIPS
SEDs); it also modifies the FIPS-compliance AK (FIPS AK) of FIPS SEDs. The
current data AK and FIPS AK of the device are required to effect changes
to the respective AKs and FIPS compliance. The current and new AKs must be
available from the key servers or onboard key management.

The command releases the cluster shell after launching the operation.
Monitor the output of the xref:{relative_path}storage-encryption-disk-
show-status.html[storage encryption disk show-status] command for command
completion.

NOTE: To properly protect data at rest on a FIPS SED and place it into
compliance with its FIPS certification requirements, set both the Data and
FIPS AKs to a value other than the device's default key; depending on the
device type, the default may be manufacture secure ID (MSID), indicated by
a key ID with the special value ``_0x0_`` , or a null key represented by a
blank key ID. Verify the key IDs by using the xref:{relative_path}storage-
```

```
encryption-disk-show.html[storage encryption disk show] and
xref:{relative_path}storage-encryption-disk-show.html[storage encryption
disk show -fips] commands.



== Parameters

``-disk <disk path name>`` - Disk Name::
This parameter specifies the name of the SED or FIPS SED that you want to
modify.

{ ``[-data-key-id <text>]`` - Key ID of the New Data Authentication Key::
This parameter specifies the key ID associated with the data AK that you
want the SED to use for future authentications. When the provided key ID
is the MSID, data at rest on the SED is not protected from unauthorized
access. Setting this parameter to a non-MSID value automatically engages
the power-on-lock protections of the device, so that when the device is
power-cycled, the system must authenticate with the device using the AK to
reenable I/O operations. You cannot specify the null default key; use MSID
instead.

| ``[-fips-key-id <text>]`` - Key ID of the New Authentication Key for
FIPS Compliance }::
This parameter specifies the key ID associated with the FIPS AK that you
want the FIPS SED to apply to SED credentials other than the one that
protects the data. When the value is not the MSID, these credentials are
changed to the indicated AK, and other security-related items are set to
conform to the FIPS certification requirements ("FIPS compliance mode") of
the device. You may set the `-fips-key-id` to any one of the key IDs known
to the system. The FIPS key ID may, but does not have to, be the same as
the data key ID parameter. Setting `-fips-key-id` to the MSID key ID value
disables FIPS compliance mode and restores the FIPS-related authorities
and other components as required (other than data) to their default
settings. A nonMSID FIPS-compliance key may be applied only to a FIPS SED.

== Examples

The following command changes both the AK and the power-cycle protection
to values that protect the data at rest on the disk. Note that the `-data-
key-id` and `-fips-key-id` parameters require one of the key IDs that
appear in the output of the `security key-manager query` command.
```

cluster1::> storage encryption disk modify -data-key-id 6A1E21D8000000000100000000000000F5A1EB48EF26FD6A8E76549C019F2350 -disk 2.10.*

Info: Starting modify on 14 disks.
View the status of the operation by using the
storage encryption disk show-status command.

```
The following command changes the FIPS AK and sets the device into FIPS-
compliance mode. Note that the `-fips-key-id` parameter requires one of
the key IDs that appear in the output of the `security key-manager query`
command.
```

cluster1::> storage encryption disk modify -fips-key-id
6A1E21D8000000000100000000000005A1FB4EE8F62FD6D8AE6754C9019F35A 2.10.*

Info: Starting modify on 14 disks.
View the status of the operation by using the
storage encryption disk show-status command.

```
== Related Links

* xref:{relative_path}storage-encryption-disk-show-status.html[storage
encryption disk show-status]
* xref:{relative_path}storage-encryption-disk-show.html[storage encryption
disk show]




[[IDefd29e7c6b7efcece0482977e1528402]]
= storage encryption disk revert-to-original-state
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Revert a self-encrypting disk to its original, as-manufactured state

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Some self-encrypting disks (SEDs) are capable of an operation that
restores them as much as possible to their as-manufactured state. The
`storage encryption disk revert-to-original-state` command invokes this
special operation that is available only in SEDs that have the physical
secure ID (PSID) printed on their labels.
```

```
The PSID is unique to each SED, meaning the command can revert only one
SED at a time. The disk must be in a "broken" or "spare" state as shown by
the output of the xref:{relative_path}storage-disk-show.html[storage disk
show] command.

The operation in the SED accomplishes the following changes:

*  Sanitizes all data by changing the disk encryption key to a new random
value
*  Sets the data authentication key (AK) and FIPS AK to the default values
*  Resets the data locking controls
*  Resets the power-on lock state to ``_false_``
*  Initializes other vendor-unique encryption-related parameters



The command releases the cluster shell after launching the operation.
Monitor the output of the xref:{relative_path}storage-encryption-disk-
show-status.html[storage encryption disk show-status] command for command
completion.

When the operation is complete, it is possible to return the SED to
service using the xref:{relative_path}storage-disk-unfail.html[storage
disk unfail] command in ``_advanced_`` privilege mode. To do so, you might
also need to reestablish ownership of the SED using the
xref:{relative_path}storage-disk-assign.html[storage disk assign] command.

== Parameters

``-disk <disk path name>`` - Disk Name::
The name of the SED to be reverted to its as-manufactured state. See the
man page for the `storage disk modify` command for information about disk-
naming conventions.

``-psid <text>`` - Physical Secure ID::
The PSID printed on the SED label.

== Examples

The following command shows a SED being returned to its as-manufactured
state:
```

cluster1::> storage encryption disk revert-to-original-state -disk 01.10.0 -psid AC65PYF8CG45YZABUQJKM98WV2VZGRLD

```
== Related Links

* xref:{relative_path}storage-disk-show.html[storage disk show]
* xref:{relative_path}storage-encryption-disk-show-status.html[storage
encryption disk show-status]
* xref:{relative_path}storage-disk-unfail.html[storage disk unfail]
* xref:{relative_path}storage-disk-assign.html[storage disk assign]




[[ID56de4477d80b6f04a3fd2b855cdf68e1]]
= storage encryption disk sanitize
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cryptographically sanitize a self-encrypting disk

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage encryption disk sanitize` command cryptographically sanitizes
one or more self-encrypting disks (SEDs), making the existing data on the
SED impossible to retrieve. This operation employs the inherent erase
capability of SEDs to perform all of the following changes:

*  Sanitizes all data by changing the disk encryption key to a new random
value
*  Sets the data authentication key (AK) to the default AK (manufacture
secure ID/MSID or null, depending on the device type)
*  Unlocks the data band
*  Resets the power-on lock state to ``_false_``



There is no method to restore the disk encryption key to its previous
value, meaning that you cannot recover the data on the SED. Use this
command with extreme care.

The sanitize command requires you to enter a confirmation phrase before
proceeding with the operation.
```

```
The command releases the cluster shell after launching the operation.
Monitor the output of the xref:{relative_path}storage-encryption-disk-
show-status.html[storage encryption disk show-status] command for command
completion.

When the operation is complete, it is possible to return the SED to
service using the xref:{relative_path}storage-disk-unfail.html[storage
disk unfail] command in ``_advanced_`` privilege mode. To do so, you might
also need to reestablish ownership of the SED using the
xref:{relative_path}storage-disk-assign.html[storage disk assign] command.

== Parameters

``-disk <disk path name>`` - Disk Name::
This parameter specifies the name of the SEDs you want to
cryptographically sanitize. See the man page for the `storage disk modify`
command for information about disk-naming conventions.

``[-force-all-states <true>]`` - Sanitize All Matching Disks::
When this parameter is ``_false_`` or not specified, the operation
defaults to spare and broken disks only, as reported in the output of the
xref:{relative_path}storage-disk-show.html[storage disk show] command.
When you specify this parameter as ``_true_`` , it allows you to
cryptographically sanitize all matching disk names regardless of their
state, including those in active use in aggregates. This allows a quick
erasure of all system data if you use the `-disk` parameter with the
asterisk wildcard (*). If you sanitize active disks, the nodes might not
be able to continue operation, and might halt or panic.

== Examples

The following command sanitizes the disk 1.10.20:
```

cluster1::> storage encryption disk sanitize 1.10.20

Warning: This operation will cryptographically sanitize 1 spare or broken
self-encrypting disk on 1 node.
To continue, enter
sanitize disk
:sanitize disk

Info: Starting sanitize on 1 disk.
View the status of the operation using the
storage encryption disk show-status command.

cluster1::>

```
If you do not enter the correct confirmation phrase, the operation is
aborted:
```

cluster1::> storage encryption disk sanitize 1.10.2*

Warning: This operation will cryptographically sanitize 5 spare or broken
self-encrypting disks on 1 node.
To continue, enter
sanitize disk
:yes
No disks sanitized.

cluster1::>

```
The following command quickly cryptographically sanitizes all system
disks, including those in active use in aggregates and shared devices:
```

cluster1::> storage encryption disk sanitize -force-all-states -disk *

Warning: This operation will cryptographically sanitize 96
self-encrypting disks on 4 nodes.
To continue, enter
sanitize disk
:sanitize disk

Info: Starting sanitize on 96 disks.
View the status of the operation by using the
storage encryption disk show-status command.

cluster1::>

```
== Related Links

* xref:{relative_path}storage-encryption-disk-show-status.html[storage
encryption disk show-status]
* xref:{relative_path}storage-disk-unfail.html[storage disk unfail]
* xref:{relative_path}storage-disk-assign.html[storage disk assign]
* xref:{relative_path}storage-disk-show.html[storage disk show]




[[ID0986229c63e878491af987af31e8ba25]]
= storage encryption disk show-status
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status of disk encryption operation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage encryption disk show-status` command displays the results of
the latest `destroy` , `modify` , or `sanitize` operation of the `storage
encryption disk` command family. Use this command to view the progress of
these operations on self-encrypting disks (SEDs).

== 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::
If you specify this parameter, the command displays disk encryption status
for the nodes that match this parameter.

``[-is-fips-support {true|false}]`` - Node Supports FIPS Disks::
If you specify this parameter, the command displays disk encryption status
for the nodes that match this parameter (``_true_`` means the node
supports FIPS-certified self-encrypting drives).

``[-latest-op <Storage Disk Encryption Operation>]`` - Latest Operation
Requested::
If you specify this parameter, the command displays disk encryption status
for the nodes with a most recent `storage encryption disk` operation that
matches this parameter (one of `destroy` , `modify` , `revert-to-original-
state` , `sanitize` , or `unknown` ).

``[-op-start-time <MM/DD/YYYY HH:MM:SS>]`` - Operation Start Time::
Selects the nodes with operation start times that match this parameter.

``[-op-execute-time <integer>]`` - Execution Time in Seconds::
If you specify this parameter, the command displays disk encryption status
```

```
for the nodes with operation execution time that matches this parameter.
The operation may be partial or completed.

``[-disk-start-count <integer>]`` - Number of Disks Started::
If you specify this parameter, the command displays disk encryption status
for the nodes that started this number of SEDs in their latest operation.

``[-disk-done-count <integer>]`` - Number of Disks Done::
Selects the nodes that report this number of SEDs having completed the
latest operation, successfully or not.

``[-disk-success-count <integer>]`` - Number of Disks Successful::
If you specify this parameter, the command displays disk encryption status
for the nodes that report this number of SEDs that successfully completed
the latest operation. When the operation is finished, if the success count
is not the same as the started count, some additional detail is available
using the `-instance` or `-node` parameters.

``[-disk-no-key-id-count <integer>]`` - Number of Disks with Key ID Not
Found::
If you specify this parameter, the command displays disk encryption status
for the nodes that report this number of SEDs that failed the latest
operation because ONTAP could not find the Key IDs associated with the
required authentication key of the SED.

``[-disk-no-authent-count <integer>]`` - Number of Disks Not
Authenticated::
If you specify this parameter, the command displays disk encryption status
for the nodes that report this number of SEDs that failed the latest
operation because the identified Authentication Key could not authenticate
with the SED.

``[-op-sequence-count <integer>]`` - Sequence Count of Latest Operation::
If you specify this parameter, the command displays disk encryption status
for that nodes that match the value list.

== Examples

When no operation has been requested since node boot, the status for that
node is empty. If you enter a node name, the output is in the same format
as for the `-instance` parameter.
```

cluster1::> storage encryption disk show-status -node node
Node Name: node
Node Supports FIPS-certified Self-Encrypting Disks: true
Latest Operation Requested: unknown
Operation Start Time: -

Execution Time in Seconds: -
Number of Disks Started: -
Number of Disks Done: -
Number of Disks Successful: -
Number of Disks with Key ID Not Found: -
Number of Disks Not Authenticated: -

```
Once an operation begins, the status is dynamic until all devices have
completed. When disks are modified, sanitized, or destroyed, sequential
executions of `storage encryption disk show-status` appear as in this
example that shows the progress of a modify operation on three SEDs on
each node of a two-node cluster:
```

cluster1::> storage encryption disk show-status
SED Latest Start Execution Disks Disks Disk
Node Support Request Timestamp Time (sec) Begun Done Successful
------- ------- -------- ---------- ---------- ------ ------ ----------
node true modify 9/22/2014 13:58:53 4 3 0 0
node1 true modify 9/22/2014 13:58:53 4 3 0 0

cluster1::> storage encryption disk show-status
SED Latest Start Execution Disks Disks Disk
Node Support Request Timestamp Time (sec) Begun Done Successful
------- ------- -------- ---------- ---------- ------ ------ ----------
node true modify 9/22/2014 13:58:53 7 3 3 3
node1 true modify 9/22/2014 13:58:53 7 3 3 3

```
[[IDc6246e6c71b29267782e17ec2ddaed18]]
= storage encryption disk show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display self-encrypting disk attributes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage encryption disk show` command displays information about
encrypting drives. When no parameters are specified, the command displays
the following information about all encrypting drives:

* Disk name
* The protection mode of the device
```

```
* The key ID associated with the data authentication key ("data AK")



In MetroCluster systems, the information is valid from the cluster that
owns the drive, or from the DR cluster when in switchover mode. If
information is not available, perform the `show` command from the cluster
partner.

You can use the following parameters together with the `-disk` parameter
to narrow the selection of displayed drives or the information displayed
about them.

== 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.

| ``[-fips  ]``::
If you specify this parameter, the command displays the key ID associated
with the FIPS-compliance authentication key ("FIPS AK") instead of the
data key ID.

| ``[-instance  ]`` }::
If you specify this parameter, the command displays detailed disk
information about all disks, or only those specified by a -disk parameter.

``[-disk <disk path name>]`` - Disk Name::
If you specify this parameter, the command displays information about the
specified disks. If you specify a single disk path name, the output is the
same as when you use the -instance parameter. See the man page for the
`storage disk modify` command for information about disk-naming
conventions. Default is all self-encrypting disks.

``[-container-name <text>]`` - Container Name::
This parameter specifies the container name associated with an encrypting
drive. If you specify an aggregate name or other container name, only the
encrypting drives in that container are displayed. See the man page for
the xref:{relative_path}storage-disk-show.html[storage disk show] command
for a description of the container name. Use the
xref:{relative_path}storage-aggregate-show-status.html[storage aggregate
show-status] and xref:{relative_path}storage-disk-show.html[storage disk
show] commands to determine which aggregates the drives are in.

``[-container-type {aggregate | broken | foreign | labelmaint |
```

```
maintenance | mediator | remote | shared | spare | unassigned | unknown |
unsupported}]`` - Container Type::
This parameter specifies the container type associated with an encrypting
drive. If you specify a container type, only the drives with that
container type are displayed. See the man page for the
xref:{relative_path}storage-disk-show.html[storage disk show] command for
a description of the container type.

``[-data-key-id <text>]`` - Key ID of the Current Data Authentication
Key::
This parameter specifies the key ID associated with the data AK that the
encrypting drive requires for authentication with its data-protection
authorities. The special key ID ``_0x0_`` indicates that the current data
AK of the drive is the default manufacture secure ID (MSID) that is not
secret. Some devices employ an initial null default AK that appears as a
blank data-key-id; you cannot specify a null data-key-id value. To
properly protect data at rest on the device, modify the data AK using a
key ID that is not a default value (MSID or null). When you modify the
data AK with a non-MSID key ID, the system automatically sets the device's
power-on lock enable control so that authentication with the data AK is
required after a device power-cycle. Use xref:{relative_path}storage-
encryption-disk-modify.html[storage encryption disk modify]`-data-key-
id```_key-id_`` to protect the data. Use xref:{relative_path}storage-
encryption-disk-modify.html[storage encryption disk modify]`-fips-key-
id```_key-id_`` to place the drives into FIPS-compliance mode.

``[-fips-key-id <text>]`` - Key ID of the Current FIPS Authentication
Key::
This parameter specifies the key ID associated with the FIPS
authentication key ("FIPS AK") that the system must use to authenticate
with FIPS-compliance authorities in FIPS-certified drives. This parameter
may not be set to a non-MSID value in drives that are not FIPS-certified.

``[-is-power-on-lock-enabled {true|false}]`` - Is Power-On Lock Protection
Enabled?::
This parameter specifies the state of the control that determines whether
the encrypting drive requires authentication with the data AK after a
power-cycle. The system enables this control parameter automatically when
you use the xref:{relative_path}storage-encryption-disk-
modify.html[storage encryption disk modify]`-data-key-id` command to set
the data AK to a value other than the default AK. Data is protected only
when this parameter is ``_true_`` and the data AK is not a default.
Compare with the values of the `-protection-mode` parameter below.

``[-protection-mode <text>]`` - Mode of SED Data and FIPS-Compliance
Protection::
```

```
The protection mode that the drive is in:
+
  * open - data is unprotected; drive is not in FIPS-compliance mode
  * data - data is protected; drive is not in FIPS-complance mode
  * part - data is unprotected; drive is otherwise in FIPS-compliance mode
  * full - data is protected; drive is in FIPS-compliance mode
  * miss - protection mode information is not available



``[-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS  | SSD | VMDISK |
SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Disk Type::
This parameter selects the drive type to include in the output.

``[-control-standard <text>]`` - Control Standard::
This parameter specifies the industry standard for control of encrypting
drives that the drive implements.

``[-compliance-standard <text>]`` - Compliance Standard::
This parameter specifies the industry compliance standard, if any, that
the drive is certified as adhering to.

``[-overall-security <text>]`` - Overall Security::
This parameter specifies the drive's certified security level as defined
in the compliance-standard, if the drive is certified to a compliance
standard.

``[-range-start <integer>]`` - TCG Data Range Start Sector::
This parameter specifies the TCG data range locking start sector.

``[-range-length <integer>]`` - TCG Data Range Length in Sectors::
This parameter specifies the TCG data range locking length in sectors.

``[-rdwr-lock-enabled <text>]`` - TCG Read/Write Lock Enabled::
This parameter indicates whether the drive will lock if it is power-cycled
when lock is enabled for read/write.

``[-rdwr-locked <text>]`` - TCG Read/Write Locked::
This parameter specifies if the drive is locked for read/write.

``[-active-key <text>]`` - TCG Active Key Type::
This parameter shows the TCG active media encryption key type used to
encrypt and decrypt the data.

== Examples

The following command displays information about all encrypting drives:
```

```
cluster1::> storage encryption disk show
Disk Mode Data Key ID
------- ---- ----------------------------------------------------------------
0.0.0 open 0x0
0.0.1 part 0x0
0.0.2 data 0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
1.10.0 open 0A53ED2A0000000001000000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
1.10.1 part 0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
1.10.2 full 0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
```

```
Note in the example that only disk 1.10.2 is fully protected with FIPS
mode, power-on-lock enable, and an AK that is not the default MSID or a
null key.

The following command displays information about the protection mode and
FIPS key ID for all encrypting drives:
```

```
cluster1::> storage encryption disk show -fips
Disk Mode FIPS-Compliance Key ID
------- ---- ----------------------------------------------------------------
0.0.0 open 0x0
0.0.1 part 0A53ED2A0000000001000000000000000C1B27AD3F0DB8891375AED2F34D0BBED
0.0.2 data 0x0
1.10.0 open 0A53ED2A0000000001000000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
1.10.1 part 0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
1.10.2 full 0A9C9CFC0000000001000000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
```

```
Note again that only disk 1.10.2 is fully protected with FIPS-compliance
mode set, power-on-lock enabled, and a data AK that is not the default
MSID or a null key.

The following command displays the individual fields for disk 1.10.2:
```

```
cluster1::> storage encryption disk show -disk 1.10.2
Disk Name: 1.10.2
Container Name: aggr0
Container Type: shared
Is Drive FIPS-certified?: true
Key ID of the Current Data Authentication Key:
0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
Key ID of the Current FIPS Authentication Key:
0A9C9CFC0000000001000000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
Is Power-On Lock Protection Enabled?: true
Mode of Data and FIPS-Compliance Protection: full
Drive Type: SSD
Control Standard: TCG Enterprise
Compliance Standard: FIPS 140-2
Overall Security: Level 2
```

```
== Related Links

* xref:{relative_path}storage-disk-show.html[storage disk show]
* xref:{relative_path}storage-aggregate-show-status.html[storage aggregate
show-status]
* xref:{relative_path}storage-encryption-disk-modify.html[storage
encryption disk modify]




:leveloffset: -1


= storage errors commands

:leveloffset: +1


[[IDdaac7e3b35facfd4acc6869e8bb520ab]]
= storage errors show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display storage configuration errors.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage errors show` command displays configuration errors with back
end storage arrays.

== 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.
```

```
``[-uid <text>]`` - UID::
Selects the disks that match this parameter value.

``[-array-name <array name>]`` - Array Name::
Selects the disks that have the specified name for the storage array that
is connected to the cluster.

``[-node {<nodename>|local}]`` - Controller Name::
Selects the disks that match this parameter value.

``[-disk <disk path name>]`` - Disk::
Selects the disks that match this parameter value.

``[-serial-number <text>]`` - Serial Number::
Selects the disks that match this parameter value.

``[-error-id <integer>,...]`` - Error ID::
Selects the disks with error-id values that match this parameter value.

``[-error-type
{onepath|onedomain|control|foreign|toobig|toosmall|invalidblocksize|target
asymmap|deviceassymmap|failovermisconfig|unknown|netapp|fwdownrev|qualfail
|diskfail|notallflashdisk}]`` - Error Type::
Selects the disks with error types values that match this parameter value.

== Examples

The following example displays configuration errors seen in the system:
```

cluster1::> storage errors show
Disk: vnv3070f20b:vnci9124s54:1-24.126L23
---
vnci9124s54:1-24.126L23 (600a0b800019e999000036b24bac3983): This array LUN reports an invalid block size and is not usable. Only a block size of 512 is supported.

```
:leveloffset: -1


= storage failover commands

:leveloffset: +1


[[IDd6ed5143320bd56a9265619dea0e0d1e]]
= storage failover check-takeover
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status for all takeover options

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover check-takeover ` command displays information about
whether takeover is possible based on verification checks triggered by the
`-check-only ` parameter used with the xref:{relative_path}storage-
failover-takeover.html[storage failover takeover ] command.

== 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>|local}]`` - Node::
Name of node command is run from

``[-operation {takeover|giveback}]`` - Operation::
Failover Operation Type

``[-option <takeover option>]`` - Takeover Option::
Takeover Option

``[-operation-status {true|false}]`` - Operation Status::
Status of whether defined operation would succeed or not

``[-reason <text>,...]`` - Reasons for Failure::
Reason for pre-takeover failure

== Examples

The following is an example of the `storage failover check-takeover`
```

```
command:
```

cluster1::*> storage failover check-takeover
Node Operation Option Operation-Status Reasons for Failure
--- --- --- --- ---
Node 1 Takeover normal false Takeover request failed due to giveback in progress

```
== Related Links

* xref:{relative_path}storage-failover-takeover.html[storage failover
takeover]




[[ID73413db5a1a15cee5acaf9320010958e]]
= storage failover giveback
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Return failed-over storage to its home node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover giveback` command returns storage that has failed
over to a node's partner back to the home node. This operation fails if
other resource-intensive operations (for instance, system dumps) are
running and make the giveback operation potentially dangerous or
disruptive. It could take anywhere from a few seconds to a couple of
minutes to relocate an SFO aggregate. The time includes both client outage
and non-outage portions. Some options are available only at the advanced
privilege level and higher. Run the xref:{relative_path}storage-failover-
show-giveback.html[storage failover show-giveback] command to check the
status of giveback operations.

NOTE:

* If the system ID of the partner has changed while the node is in
takeover mode, the `storage failover giveback` command updates the
ownership of the partner's disks to the new system ID while giving back.
```

```
* If the giveback operation fails due to the operation being vetoed by a
subsystem, check the syslog or EMS output for a subsystem-specific reason
for the abort. The corrective action is subsystem-specific and is detailed
in the corrective action portion of the message. Follow the corrective
action specified by the subsystem and then reissue the `storage failover
giveback` command. If you cannot perform the corrective action, then use
the `override-vetoes` option in the `storage failover giveback` command to
force the giveback.
* If the giveback operation fails because the node cannot communicate with
its partner, check the EMS output for the corrective action. Follow the
corrective action and then reissue the `storage failover giveback`
command. If you cannot perform the corrective action, then use the `-
require-partner-waiting false` option in the `storage failover giveback`
command to force the giveback. This parameter is available only at the
advanced privilege level and higher.
* If the node does not receive notification that the partner has brought
online the given-back aggregate and its volumes, the
xref:{relative_path}storage-failover-show-giveback.html[storage failover
show-giveback] command displays the giveback status for the aggregate as
failed. A possible reason for this failure is that the partner is
overloaded and slow in bringing the aggregate online. Run the
xref:{relative_path}storage-aggregate-show.html[storage aggregate show]
command to verify that the aggregate and its volumes are online on the
partner node. The node will not attempt the giveback operation for
remaining aggregates. To force the giveback, use the `-require-partner-
waiting false` option in the `storage failover giveback` command. This
parameter is available only at the advanced privilege level and higher.




== Parameters

{ ``-ofnode {<nodename>|local}`` - Node to which Control is Givenback::
Specifies the node whose storage is currently taken over by its partner
and will be given back by the giveback operation.

| ``-fromnode {<nodename>|local}`` - Node Initiating Giveback }::
Specifies the node that currently holds the storage that is to be returned
to the partner node.

``[-require-partner-waiting {true|false}]`` - Require Partner in Waiting
[.advaced-priv]#(privilege: advanced)#::
 If this optional parameter is used and set to false, the storage is given
back regardless of whether the partner node is available to take back the
```

```
storage or not. If this parameter is used and set to true, the storage
will not be given back if the partner node is not available to take back
the storage. If this parameter is not used, the behavior defaults to the
setting of the `-check-partner` option set with the
xref:{relative_path}storage-failover-modify.html[storage failover modify]
command.

``[-override-vetoes <true>]`` - Override All Vetoes::
 If this optional parameter is used, the system overrides veto votes
during a giveback operation. If this parameter is not used, the system
does not proceed with a giveback if it is vetoed. This parameter, if used,
can only be set to true.

``[-only-cfo-aggregates <true>]`` - Giveback Only CFO Aggregates::
If this optional parameter is used, giveback of only the CFO aggregates
(root aggregate and CFO style data aggregates) will be attempted. If this
parameter is not used, giveback of all the aggregates (CFO and SFO
aggregates) will be attempted. This parameter, if used, can only be set to
true.

== Examples

The following example gives back storage that is currently held by a node
named node1. The partner must be available for the giveback operation to
occur.
```

node::> storage failover giveback -fromnode node1

```
The following example gives back only the CFO aggregates to a node named
node2 (the aggregates are currently held by a node named node1). The
partner must be available for the giveback operation to occur, and the
veto-giveback process can be overridden.
```

node::> storage failover giveback -ofnode node2
-override-vetoes true -only-cfo-aggregates true

```
== Related Links

* xref:{relative_path}storage-failover-show-giveback.html[storage failover
show-giveback]
* xref:{relative_path}storage-aggregate-show.html[storage aggregate show]
* xref:{relative_path}storage-failover-modify.html[storage failover
modify]
```

```
[[ID3df2deedcc4dbeb2b33ccdea1ed30b6b]]
= storage failover modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify storage failover attributes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover modify` command changes the storage-failover options
for a node. Some options are available only at the advanced privilege
level and higher.

== Parameters

``-node {<nodename>|local}`` - Node::
This specifies the node whose storage-failover options are to be modified.

{ ``[-enabled {true|false}]`` - Takeover Enabled::
This optionally specifies whether storage failover is enabled. The default
setting is `true` .

| ``[-mode {ha|non_ha}]`` - HA Mode }::
This specifies whether the node is set up in high-availability mode or
stand-alone mode. If the node is a member of a high-availability
configuration, set the value to ``_ha_`` . If the node is stand-alone, set
the value to ``_non_ha_`` . Before setting the HA mode, you must complete
the platform dependent steps to set up the system in a stand-alone or HA
configuration as shown in the documentation for your platform.

``[-auto-giveback {true|false}]`` - Auto Giveback Enabled::
This optionally specifies whether automatic giveback operations are
enabled. An automatic giveback operation is invoked when one node of a
failover pair is in takeover mode and the failed node is repaired and
restarts. When the repaired node boots, the node in takeover mode detects
this and initiates a giveback operation. The default setting is ``_true_``
. This parameter is not applicable if takeover was because of a disruption
in the partner's operation. For those cases, use the `-auto-giveback-
after-panic` parameter, instead.
```

```
``[-detection-time <integer>]`` - Takeover Detection Time (secs)::
This optionally specifies the amount of time, in seconds, that a node
remains unresponsive before its partner initiates a takeover operation.
Possible values range from 10 to 180 seconds. The default setting is 15
seconds.

``[-onfailure {true|false}]`` - Takeover on Failure Enabled [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether the node automatically takes over for
its partner node if the partner node fails. The default setting is `true`
.This parameter is available only at the advanced privilege level and
higher.

``[-onpanic {true|false}]`` - Takeover on Panic Enabled::
This optionally specifies whether the node automatically takes over for
its partner node if the partner node panics. The default setting is `true`
. Changing this parameter on one node automatically makes the same change
on its partner node.

``[-onshort-uptime {true|false}]`` - Takeover on Short Uptime Enabled
[.advaced-priv]#(privilege: advanced)#::
This optionally specifies whether the node takes over for its partner node
if the partner node fails within 60 seconds of starting up; the time
period is modifiable by using the `-short-uptime` parameter. The default
setting is `true` . This parameter is available only at the advanced
privilege level and higher.

``[-short-uptime <integer>]`` - Short Uptime (secs) [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies the time period used by the `-onshort-uptime`
parameter. The default setting is 60 seconds. This parameter is available
only at the advanced privilege level and higher.

``[-attempts <integer>]`` - Number of Giveback Attempts [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies the number of times the node attempts an
automatic giveback operation within 60 minutes; the time period is
modifiable by using the `-attempts-time` parameter. The default setting is
2 attempts. This parameter is available only at the advanced privilege
level and higher.

``[-attempts-time <integer>]`` - Giveback Attempts Period (minutes)
[.advaced-priv]#(privilege: advanced)#::
This optionally specifies the time period used by the `-attempts`
parameter. The default setting is 60 minutes. This parameter is available
```

```
only at the advanced privilege level and higher.

``[-propagate {true|false}]`` - Propagate Status via Mailbox [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether storage-failover status is communicated
via mailbox disks. The default setting is `true` . This parameter is
available only at the advanced privilege level and higher.

``[-read-interval <integer>]`` - Node Status Read Interval (secs)
[.advaced-priv]#(privilege: advanced)#::
This optionally specifies, in seconds, how frequently the node reads its
partner node's status from the mailbox disks. The default setting is 5
seconds. This parameter is available only at the advanced privilege level
and higher.

``[-write-interval <integer>]`` - Node Status Write Interval (secs)
[.advaced-priv]#(privilege: advanced)#::
This optionally specifies, in seconds, how frequently the node writes its
status to the mailbox disks. The default setting is 5 seconds. This
parameter is available only at the advanced privilege level and higher.

``[-onreboot {true|false}]`` - Takeover on Reboot Enabled::
This optionally specifies whether the node automatically takes over for
its partner if the partner reboots. The default setting is `true` .
Takeover can occur if the partner exceeds the expected time to reboot even
when this option is set to false. The expected time to reboot is different
for different platforms. The minimum expected time to reboot is 180
seconds. The `-inhibit-takeover` option of the xref:{relative_path}system-
node-reboot.html[system node reboot] command overrides this option: if a
node is rebooted with `-inhibit-takeover` set to `true` then takeover does
not occur, even if the `takeover on reboot` option is true. If a node does
takeover due to the partner rebooting, then it will automatically giveback
after the reboot, even if the `-auto-giveback` option is set to `false` .
This is non-persistent behavior: if the node does takeover due to partner
reboot and then itself reboots (prior to giveback) then it will not
automatically giveback if the `-auto-giveback` option is set to `false` .

``[-delay-seconds <integer>]`` - Delay Before Auto Giveback (secs)::
This optionally specifies the minimum time that a node will stay in
takeover state prior to performing an automatic giveback. If the taken
over node recovers quickly (for example, if the takeover was due to a
reboot), by delaying the giveback for a few minutes the outage during the
takeover and giveback can be reduced to two short outages instead of one
longer one. The allowed range is 0 to 600, inclusive. The default setting
is 600 seconds. This option affects all types of auto-giveback. This
parameter is available only at the advanced privilege level and higher.
```

```
+
NOTE: This delay does not affect manual giveback.
+


``[-hwassist {true|false}]`` - Hardware Assist Enabled::
This optionally specifies whether the hardware assist feature is enabled.
If set to `true` this feature helps in fast takeover detection times in
certain cases.

``[-hwassist-partner-ip <IP Address>]`` - Partner's Hwassist IP::
This optionally specifies the Ip address on which the partner node
receives hardware assist alerts. For the hardware assist feature to be
active,the value of this option should be equal to partner's node
management Ip address.

``[-hwassist-partner-port <integer>]`` - Partner's Hwassist Port::
This optionally specifies the port number on which partner node listens to
hardware assist alerts.It is recommended to have this value to be between
4000-4500. The default value is 4444.

``[-hwassist-health-check-interval <integer>]`` - Hwassist Health Check
Interval (secs)::
This optionally specifies,in seconds, how frequently the hardware assist
hardware on a node sends a heartbeart to it's partner.The default value is
180.

``[-hwassist-retry-count <integer>]`` - Hwassist Retry Count::
This optionally specifies the number of times we repeat sending an
hardware assist alert. The default value is 2.
+
NOTE: Modification of hardware assist parameters such as hardware assist
IP, port is allowed only on certain platforms. Refer to Knowledge Base
article link:https://kb.netapp.com/on-prem/ontap/OHW/OHW-
KBs/How_to_change_HWassist_IP_address[How to change hardware assist IP
address?]
+


``[-auto-giveback-after-panic {true|false}]`` - Auto Giveback After
Takeover On Panic::
This optionally specifies whether a node should attempt automatic giveback
operations if takeover was because of a disruption in the partner's
operation. An automatic giveback operation is invoked when one node of a
failover pair is in takeover mode and the failed node is repaired and
restarts. When the repaired node boots, the node in takeover mode detects
```

```
this and initiates a giveback operation automatically. The default setting
is ``_true_`` .
+
NOTE: This parameter is independent of the `-auto-giveback` parameter. If
this parameter is enabled and the takeover is due to disruption in the
partner's operation, giveback will be initiated even if `-auto-giveback`
parameter is ``_false_`` .
+


``[-aggregate-migration-timeout <integer>]`` - Aggregate Migration Timeout
(secs) [.advaced-priv]#(privilege: advanced)#::
This optionally specifies the amount of time, in seconds, the source node
has to wait for the destination node to complete the aggregate migration
before declaring the migration as failed. The default setting is 120
seconds.

== Examples

The following example enables the storage-failover service on a node named
node0:
```

node::> storage failover modify -node node0 -enabled true

```
The following examples enable storage-failover takeover on a short uptime
of 30 seconds on a node named node0:
```

node::*> storage failover modify -node node0 -onshort-uptime true -short-uptime 30

```
== Related Links

* xref:{relative_path}system-node-reboot.html[system node reboot]




[[ID4bcb2765f60067ab73769bb47e1079d8]]
= storage failover show-giveback
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display giveback status
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover show-giveback` command displays information about
the giveback status of high-availability (HA) partner aggregates.  The
command displays the following information when no parameters are
specified:

* Node name
* Partner aggregate name
* Giveback Status

 You can specify additional parameters to display only the information
that matches those parameters. For example, to display information only
about a particular aggregate, run the command with the `-aggregate
aggregate_name ` parameter.

== 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>|local}]`` - Node::
If this parameter is used, the command displays information about the
giveback status of the aggregates belonging to the HA partner of the
specified node.

``[-aggregate <text>]`` - Aggregate::
If this parameter is used, the command displays information about the
giveback status of the specified aggregate.

``[-giveback-status <text>,...]`` - Aggregates Giveback State::
If this parameter is used, the command displays information about the
aggregates with the specified giveback status.

``[-destination <text>]`` - Destination for Giveback::
If this parameter is used, the command displays information about the
giveback status of the aggregates whose destination after the giveback is
```

```
the specified node.

== Examples

The following example displays information about giveback status on all
nodes:
```

node::> storage failover show-giveback
Partner
Node Aggregate Giveback Status
---------- ---------- ----------
node0
- No aggregates to give back
node1
- No aggregates to give back
node2
- No aggregates to give back
node3
- No aggregates to give back
4 entries were displayed.

```
[[ID8a1625b5ee52ef8b4d9754ee34948d1e]]
= storage failover show-takeover
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display takeover status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover show-takeover` command displays information about
the takeover status of nodes in a cluster. The command also displays the
takeover status of aggregates being taken over. During each phase of
takeover, the takeover node and the node being taken over display their
takeover status and the status of the aggregates being taken over.  The
command displays the following information when no parameters are
specified:

* Node name
* Node takeover status - This contains a descriptive information about the
phase of takeover.
```

```
* Aggregate
* Aggregate takeover status - This contains the following information:

  *  Takeover status of the aggregate, such as "Done", "Failed", "In
progress" and "Not attempted yet".
  *  Reason for an aggregate takeover failure.
  *  Corrective action, in case of an aggregate takeover failure.



 You can specify additional parameters to display only the information
that matches those parameters. For example, to display information only
about a particular node, run the command with the `-node```_node_name_``
parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
If this parameter is specified, the command displays the specified fields
for all nodes, in column style output.

| ``[-instance  ]`` }::
If this parameter is specified, the command displays the same detailed
information as for the `-node` parameter, but for all nodes.

``[-node {<nodename>|local}]`` - Node Name::
If this parameter is specified, the command displays information about the
takeover status of the specified node, and the takeover status of the
aggregates being taken over.

``[-node-takeover-status <text>]`` - Node's Takeover Status::
If this parameter is specified, the command displays information about the
takeover status of the nodes with the specified node-takeover-status. The
command also displays the takeover status of the aggregates belonging to
the node being taken over.

``[-aggregate <text>]`` - Aggregate Being Taken Over::
If this parameter is specified, the command displays information about the
takeover status of the specified aggregate, and the takeover status of the
nodes containing the specified aggregate.

``[-aggregate-takeover-status <text>]`` - Aggregate's Takeover Status::
If this parameter is specified, the command displays information about the
takeover status of the aggregates with the specified aggregate takeover
status, and the takeover status of the nodes containing those aggregates.

== Examples
```

```
The following example shows the takeover status of two nodes, nodeA and
nodeB, in an High Availability (HA) pair, when both are in normal mode;
neither node has taken over its HA partner. In this case, there is no
takeover status for the aggregates.
```

cluster1::> storage failover show-takeover
Node Node Status Aggregate Takeover Status
---------- -------------------- -------------- ------------------------------
nodeA Takeover not
attempted.
- -
nodeB Takeover not
attempted.
- -

```
The following example shows the takeover status of two nodes, nodeA and
nodeB, in an HA pair, when nodeA is in the SFO phase of an optimized
takeover of nodeB. In this case, nodeA does not have information about the
takeover status of nodeB's aggregates.
```

cluster1::> storage failover show-takeover
Node Node Status Aggregate Takeover Status
---------- -------------------- -------------- ------------------------------
nodeA Optimized takeover
of partner in
progress. Partner,
("nodeB"), is
relocating its SFO
aggregates. Run the
command "storage
failover
show-takeover -node
nodeB" to display the
relocation status of
the partner.
- -
nodeB Being taken over.
aggr1 In progress, Module: backup.
aggr2 Not attempted yet
CFO aggregates Not attempted yet.

```
The following example shows the takeover status of two nodes, nodeA and
nodeB, in an HA pair, when nodeA has completed the SFO phase of an
optimized takeover of nodeB (but has not completed the CFO phase of the
optimized takeover). In this case, nodeA has information about the
takeover status of nodeB's aggregates.
```

cluster1::> storage failover show-takeover
ode Node Status Aggregate Takeover Status
--------- --------------------- -------------- -------------------------------
odeA Partner has
relocated its
aggregates. Takeover
in progress.
aggr1 Done
aggr2 Done
CFO aggregates In progress.
odeB Relocated aggregates
to partner. Waiting
for partner to
takeover.
aggr1 Done
aggr2 Done
CFO aggregates Not attempted yet.

```
The following example shows the takeover status of two nodes, nodeA and
nodeB, in an HA pair, when nodeA has completed the SFO and CFO phases of
an optimized takeover of nodeB. In this case, nodeA has information about
the takeover status of nodeB's aggregates. Since nodeB is not operational,
an Remote Procedure Call(RPC) error is indicated in the command output.
```

cluster1::> storage failover show-takeover
ode Node Status Aggregate Takeover Status
--------- --------------------- -------------- -------------------------------
odeA Partner has
relocated its
aggregates. In
takeover.
aggr1 Done
aggr2 Done
CFO aggregates Done.
arning: Unable to list entries on node nodeB. RPC: Port mapper failure - RPC:
Timed out

```
The following example shows the takeover status of two nodes, nodeA and
nodeB, in an HA pair, when nodeA has aborted the SFO phase of an optimized
takeover of nodeB. In this case, nodeA does not have information about the
takeover status of nodeB's aggregates.
```

cluster1::> storage failover show-takeover
ode Node Status Aggregate Takeover Status
--------- --------------------- -------------- -------------------------------
odeA Optimized takeover
of partner aborted.
Run the command
"storage failover

show-takeover -node
nodeB" to display the
relocation status of
the partner.
- -
odeB Optimized takeover
by partner aborted.
aggr1 Failed: Destination node did
not online the aggregate on
time. To takeover the
remaining aggregates, run the
"storage failover takeover
-ofnode nodeB
-bypass-optimization true"
command. To giveback the
relocated aggregates, run the
"storage failover giveback
-ofnode nodeB" command.
aggr2 Not attempted yet
CFO aggregates Not attempted yet.

```
[[ID9abd5c391ad492cca7fa38d15f87777b]]
= storage failover show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display storage failover status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover show` command displays information about storage-
failover configurations. By default, the command displays the following
information:

* Node name.
* Partner node name.
* Whether storage failover is possible.
* The current state of storage failover. If the takeover is disabled the
appropriate reason would be displayed.

To display detailed information about storage failover on a specific node,
run the command with the -node parameter. The detailed view adds the
following information:
```

```
* Node NVRAM ID.
* Partner NVRAM ID.
* Whether storage failover is enabled.
* Whether the storage-failover interconnect is available.
* Status of individual storage-failover interconnect links.
* Type and vendor of the storage-failover interconnect.
* Partner State
* Status codes from the takeover-by-partner process. Possible values
include:

  * NVRAM_DOWN
  * OPERATOR_DISABLE_NVRAM
  * PARTNER_RESET
  * FM_TAKEOVER
  * NVRAM_MISMATCH
  * OPERATOR_DENY
  * CLUSTER_DISABLE
  * VERSION
  * SHELF_HOT
  * REVERT_IN_PROGRESS
  * HALT_NOTKOVER
  * TAKEOVER_ON_PANIC


* Reasons why takeover is not possible, if applicable. Possible values
include:

  * NOT_INIT
  * DISABLED
  * DEGRADED
  * MBX_UNKNOWN
  * FM_VERSION
  * PARTNER_DISABLED
  * OPERATOR_DENY
  * NVRAM_MISMATCH
  * VERSION
  * IC_ERROR
  * BOOTING
  * SHELF_HOT
  * PARTNER_REVERT_IN_PROGRESS
  * LOCAL_REVERT_IN_PROGRESS
  * PARTNER_TAKEOVER
  * LOCAL_TAKEOVER
  * HALT_NOTKOVER
  * LOG_UNSYNC
```

```
  * UNKNOWN
  * WAITING_FOR_PARTNER
  * LOW_MEMORY
  * HALTING
  * MBX_UNCERTAIN
  * NO_AUTO_TKOVER


* Time until takeover, in seconds.
* Time until auto giveback, in seconds.
* Delay for auto giveback, in seconds.
* List of local mailbox disks.
* List of partner mailbox disks.
* Whether operator-initiated planned takeover will be optimized for
performance by relocating SFO (non-root) aggregates serially to the
partner prior to takeover.

You can specify additional parameters to select the displayed information.
For example, to display information only about storage-failover
configurations whose interconnect is down, run the command with`
-interconnect-up false` .

== 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.

| ``[-advanced  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information:
+
  * Node name
  * Whether kill messages are issued during a takeover operation
  * Whether the node controls its partner's storage aggregates
  * The time when firmware notification was received
  * The time when booting notification was received
  * The time at which the last takeover or giveback operation occurred, in
microseconds
  * The number of times the failover log was unsynchronized



| ``[-iotime  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information:
+
```

```
  * Node name
  * Primary normal I/O time
  * Primary transition I/O time
  * Backup normal I/O time
  * Backup transition I/O time



| ``[-mailbox-status  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information:
+
  * Node name
  * Primary mailbox status
  * Backup mailbox status



| ``[-more-options  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information:
+
  * Node name
  * Whether takeover on short uptime is enabled
  * Short uptime, in seconds
  * Number of giveback attempts
  * Interval of giveback attempts, in minutes
  * Whether the primary mailbox is online
  * Mailbox status read interval, in seconds
  * Mailbox status write interval, in seconds



| ``[-options  ]``::
Displays the following information:
+
  * Node name
  * Whether automatic giveback operations are enabled
  * Whether long-running operations are terminated when an automatic
giveback operation is initiated
  * Whether the node checks its partner's readiness before initiating a
giveback operation
  * The time, in seconds, that the node remains unresponsive before its
partner initiates a takeover operation
  * Whether the node automatically takes over for its partner if the
partner fails
  * Whether the node automatically takes over for its partner if the
partner panics
```

```
  * Whether the node automatically takes over for its partner if the
partner reboots
  * whether Hardware Assisted Takeover is enabled
  * Ip address on which the partner node listens to the Hardware Assist
alerts
  * Port number on which the partner node listens to the Hardware Assist
alerts
  * Whether operator-initiated planned takeover will be optimized for
performance by relocating SFO (non-root) aggregates serially to the
partner prior to takeover

+
If this parameter is specified when the privilege level is set to advanced
or higher, the command displays the information in the previous list and
the following additional information:
+
  * Whether the node takes over for its partner if its partner fails after
a period of time, which is listed in the following field
  * The number of seconds before the node takes over for its partner
  * The number of times the node attempts an automatic giveback operation
within a period of time
  * The number of minutes in which the automatic giveback attempts can
occur
  * Whether storage-failover status is communicated via mailbox disks
  * The interval at which the node reads its partner node's status from
the mailbox disks
  * The interval at which the node writes its status to the mailbox disks
  * '

  * The interval at which Hardware assist h/w sends a heartbeat
  * The number of times the Hardware assist alert is sent



| ``[-progress  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information:
+
  * Node name
  * Maximum resource-table index number
  * Current resource-table index number
  * Current resource-table entry



| ``[-takeover-status  ]``::
Displays the following information:
```

```
+
  * Node name
  * Partner name
  * Takeover enabled
  * Takeover possible
  * Interconnect up
  * State
  * Node NVRAM ID
  * Partner NVRAM ID
  * Reason Takeover Not Possible By Partner
  * Reason Takeover Not Possible
  * Time Until Takeover



| ``[-timeout  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information:
+
  * Node name
  * Fast timeout
  * Slow timeout
  * Mailbox timeout
  * Connection timeout
  * Operator timeout
  * Firmware timeout
  * Dump-core timeout
  * Booting timeout
  * Reboot timeout



| ``[-transit  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the following information:
+
  * Node name
  * Transit Timer Enabled
  * Transit Timeout



| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Selects the nodes whose name matches this parameter value.
```

```
``[-partner-name <text>]`` - Partner Name::
Selects the nodes that have the specified partner-name setting.

``[-nvramid <integer>]`` - Node NVRAM ID::
Selects the nodes that have the specified NVRAM ID setting.

``[-partner-nvramid <integer>]`` - Partner NVRAM ID::
Selects the nodes that have the specified partner NVRAM ID setting.

``[-enabled {true|false}]`` - Takeover Enabled::
Selects the nodes that have the specified takeover-enabled setting.

``[-mode {ha|non_ha}]`` - HA Mode::
Selects the nodes that have the specified HA-mode setting. If the value is
set to ``_ha_`` then the node is a member of a storage-failover
configuration. If it is set to ``_non-ha_`` then it is in a stand alone
configuration.

``[-possible {true|false}]`` - Takeover Possible::
Selects the nodes that have the specified failover-possible setting.

``[-reason <text>,...]`` - Reason Takeover not Possible::
Selects the nodes that have the specified reason-not-possible setting.
Possible values include:
+
  * NOT_INIT
  * DISABLED
  * DEGRADED
  * MBX_UNKNOWN
  * FM_VERSION
  * PARTNER_DISABLED
  * OPERATOR_DENY
  * NVRAM_MISMATCH
  * VERSION
  * IC_ERROR
  * BOOTING
  * SHELF_HOT
  * PARTNER_REVERT_IN_PROGRESS
  * LOCAL_REVERT_IN_PROGRESS
  * PARTNER_TAKEOVER
  * LOCAL_TAKEOVER
  * HALT_NOTKOVER
  * LOG_UNSYNC
  * UNKNOWN
  * WAITING_FOR_PARTNER
  * LOW_MEMORY
```

```
  * HALTING
  * MBX_UNCERTAIN
  * NO_AUTO_TKOVER



``[-interconnect-up {true|false}]`` - Interconnect Up::
Selects the nodes that have the specified interconnect-up setting.

``[-interconnect-links <text>]`` - Interconnect Links::
Selects the nodes that have the specified interconnect-links setting.

``[-interconnect-type <text>]`` - Interconnect Type::
Selects the nodes that have the specified interconnect-type setting.

``[-state-description <text>]`` - State Description::
Selects the nodes that have the specified state-description setting.

``[-partner-state <text>]`` - Partner State::
Selects the nodes that have the specified partner-state setting. Possible
values include:
+
  * OPERATOR COMPLETED
  * DEBUGGUER COMPLETED
  * PROGRESS COUNTER
  * I/O ERROR
  * BAD CHECKSUM
  * RESERVED
  * UNKNOWN
  * INITIALIZING
  * IN POWER-ON SELF TEST
  * BOOTING
  * BOOT FAILED
  * WAITING
  * KERNEL LOADED
  * UP
  * IN DEBUGGER
  * WAITING FOR OPERATOR INPUT
  * DUMPING CORE
  * HALTED
  * REBOOTING
  * WAITING FOR GIVEBACK (DISK RESERVATIONS)
  * WAITING FOR GIVEBACK (HA MAILBOXES)
  * DUMPING SPARECORE
  * MULTI-DISK PANIC
  * IN TAKEOVER
```

```
``[-time-until-takeover <integer>]`` - Time Until Takeover::
Selects the nodes that have the specified time-until-takeover setting.

``[-partner-reason <text>,...]`` - Reason Takeover not Possible by
Partner::
Selects the nodes that have the specified partner-reason text setting.

``[-killpackets {true|false}]`` - Issue Kill Packets [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified kill packets setting.

``[-partner-aggregates {true|false}]`` - Control Partner Aggregates
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner aggregates setting.

``[-current-index <integer>]`` - Current Progress Index [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified current-progress index setting.

``[-current-entry <text>]`` - Current Progress Entry [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified current-progress entry setting.

``[-maximum-index <integer>]`` - Maximum Progress Index [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified maximum-progress index setting.

``[-pmbox-status <text>,...]`` - Primary Mailbox Status [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified primary mailbox status setting.
Possible values include:
+
  * MBX_STATUS_NODISKS
  * MBX_STATUS_UNCERTAIN
  * MBX_STATUS_STALE
  * MBX_STATUS_CONFLICTED
  * MBX_STATUS_OLD_VERSION
  * MBX_STATUS_NOT_FOUND
  * MBX_STATUS_WRONG_STATE
  * MBX_STATUS_BACKUP



``[-bmbox-status <text>,...]`` - Backup Mailbox Status [.advaced-
priv]#(privilege: advanced)#::
```

```
Selects the nodes that have the specified backup-mailbox status setting.
See the description of the `-pmbox-status` parameter for a list of
possible values.

``[-major-seq-num-local <integer>]`` - Local Major Sequence Number
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified mailbox heartbeat major sequence
number on the local node.

``[-minor-seq-num-local <integer>]`` - Local Minor Sequence Number
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified mailbox heartbeat minor sequence
number on the local node.

``[-major-seq-num-partner <integer>]`` - Partner Major Sequence Number
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified mailbox heartbeat major sequence
number on the partner node.

``[-minor-seq-num-partner <integer>]`` - Partner Minor Sequence Number
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified mailbox heartbeat minor sequence
number on the partner node.

``[-local-mbx-node-status <Mailbox Status>]`` - Local Mailbox Node Status
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified local mailbox node status.
Possible values include:
+
  * MBX_UNKNOWN - Local node is up, mailbox uninitialized
  * MBX_TAKEOVER_DISABLED - Local node is up but takeover is disallowed
  * MBX_TAKEOVER_ENABLED - Local node is up and takeover is allowed
  * MBX_TAKEOVER_ACTIVE - Partner node has taken over
  * MBX_GIVEBACK_DONE - Giveback completed, but local node has not yet
restarted



``[-mbx-abs-time-local <integer>]`` - Local Mailbox Absolute Time
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified local mailbox channel absolute
time. This time is measured in msecs since 1/1/1970 (epoch).

``[-mbx-sk-time-local <integer>]`` - Local Mailbox Kernel Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified local mailbox channel Kernel
Time.
```

```
``[-mbx-sk-cycles-local <integer>]`` - Local Mailbox CPU Cycles [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified local mailbox channel CPU Cycle
count.

``[-ic-abs-time-local <integer>]`` - Local IC Absolute Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified local Interconnect channel
absolute time. This time is measured in msecs since 1/1/1970 (epoch).

``[-ic-sk-time-local <integer>]`` - Local IC Kernel Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified local Interconnect channel
Kernel Time.

``[-ic-sk-cycles-local <integer>]`` - Local IC CPU Cycles [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified local Interconnect channel CPU
Cycle count.

``[-partner-mbx-node-status <Mailbox Status>]`` - Partner Mailbox Node
Status [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner mailbox node status.
Possible values include:
+
  * MBX_UNKNOWN
  * MBX_TAKEOVER_DISABLED
  * MBX_TAKEOVER_ENABLED
  * MBX_TAKEOVER_ACTIVE
  * MBX_GIVEBACK_DONE



``[-mbx-abs-time-partner <integer>]`` - Partner Mailbox Absolute Time
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner mailbox channel absolute
time. This time is measured in msecs since 1/1/1970 (epoch).

``[-mbx-sk-time-partner <integer>]`` - Partner Mailbox Kernel Time
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner mailbox channel Kernel
Time.

``[-mbx-sk-cycles-partner <integer>]`` - Partner Mailbox CPU Cycles
[.advaced-priv]#(privilege: advanced)#::
```

```
Selects the nodes that have the specified partner mailbox channel CPU
Cycle count.

``[-mbx-major-seq-num-partner <integer>]`` - Partner Mailbox Major
Sequence Number [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner mailbox channel major
sequence number.

``[-mbx-minor-seq-num-partner <integer>]`` - Partner Mailbox Minor
Sequence Number [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner mailbox channel minor
sequence number.

``[-ic-abs-time-partner <integer>]`` - Partner IC Absolute Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner Interconnect channel
absolute time. This time is measured in msecs since 1/1/1970 (epoch).

``[-ic-sk-time-partner <integer>]`` - Partner IC Kernel Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner Interconnect channel
Kernel Time.

``[-ic-sk-cycles-partner <integer>]`` - Partner IC CPU Cycles [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner Interconnect channel CPU
Cycle count.

``[-ic-major-seq-num-partner <integer>]`` - Partner IC Major Sequence
Number [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner Interconnect channel
major sequence number.

``[-ic-minor-seq-num-partner <integer>]`` - Partner IC Minor Sequence
Number [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner Interconnect channel
minor sequence number.

``[-local-takeover-info <text>]`` - Local Takeover Info [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified local node takeover information.
This includes the type of negotiated failover request, or if takeover is
not possible, the reason why takeover is disabled. Possible values
include:
+
  * NOTKOVER_NVRAM_DOWN - NVRAM mirror is down
```

```
  * NOTKOVER_OPERATOR_DISABLE_NVRAM - Operator disabled
  * NOTKOVER_PARTNER_RESET - A link reset is in progress
  * NOTKOVER_FM_TAKEOVER - The failover monitor has declared takeover
  * NOTKOVER_NVRAM_MISMATCH - NVRAM sizes mismatch
  * NOTKOVER_OPERATOR_DENY - Operator denies takeover
  * NOTKOVER_CLUSTER_DISABLE - Cluster is disabled
  * NOTKOVER_VERSION - Version mismatch
  * NOTKOVER_SHELF_HOT - Disk shelf is too hot
  * NOTKOVER_REVERT_IN_PROGRESS - Revert is in progress
  * NOTKOVER_HALT_NOTKOVER - Node halted in no-takeover mode
  * TKOVER_ON_REBOOT - Enable takeover on reboot
  * TKOVER_ON_PANIC - Enabled takeover on panic
  * TKOVER_ON_STUTTER_DISABLED - Disable takeover on short uptime
  * NFO_DISK_SHELF_ENABLED - Negotiated failover for disk shelf module is
enabled
  * NFO_ISCSI_ENABLED - Negotiated failover for network interfaces module
is enabled
  * NFO_FCP_TARGET_ENABLED - Negotiated failover for fcp target module is
enabled



``[-partner-takeover-info <text>]`` - Partner Takeover Info [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner node takeover
information. This includes the type of negotiated failover request, or if
takeover is not possible, the reason why takeover is disabled. Possible
values include:
+
  * NOTKOVER_NVRAM_DOWN - NVRAM mirror is down
  * NOTKOVER_OPERATOR_DISABLE_NVRAM - Operator disabled
  * NOTKOVER_PARTNER_RESET - A link reset is in progress
  * NOTKOVER_FM_TAKEOVER - The failover monitor has declared takeover
  * NOTKOVER_NVRAM_MISMATCH - NVRAM sizes mismatch
  * NOTKOVER_OPERATOR_DENY - Operator denies takeover
  * NOTKOVER_CLUSTER_DISABLE - Cluster is disabled
  * NOTKOVER_VERSION - Version mismatch
  * NOTKOVER_SHELF_HOT - Disk shelf is too hot
  * NOTKOVER_REVERT_IN_PROGRESS - Revert is in progress
  * NOTKOVER_HALT_NOTKOVER - Node halted in no-takeover mode
  * TKOVER_ON_REBOOT - Takeover on reboot is enabled
  * TKOVER_ON_PANIC - Takeover on panic is enabled
  * TKOVER_ON_STUTTER_DISABLED - Disable takeover on short uptime
  * NFO_DISK_SHELF_ENABLED - Negotiated failover for disk shelf module is
enabled
  * NFO_ISCSI_ENABLED - Negotiated failover for network interfaces module
```

```
is enabled
  * NFO_FCP_TARGET_ENABLED - Negotiated failover for fcp target module is
enabled



``[-local-headswap-state <Headswap State>]`` - Local Head Swap State
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified local node headswap state.
Possible values are:
+
  * HEADSWAP_NONE - head swap not in progress
  * HEADSWAP_START - head swap started
  * HEADSWAP_CFO_START - CFO phase of head swap started
  * HEADSWAP_CFO_END - CFO phase of head swap completed
  * HEADSWAP_SFO_START - SFO phase of head swap started



``[-partner-headswap-state <Headswap State>]`` - Partner Head Swap State
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified partner node headswap state.
Possible values are:
+
  * HEADSWAP_NONE - head swap not in progress
  * HEADSWAP_START - head swap started
  * HEADSWAP_CFO_START - CFO phase of head swap started
  * HEADSWAP_CFO_END - CFO phase of head swap completed
  * HEADSWAP_SFO_START - SFO phase of head swap started



``[-fast-timeout <integer>]`` - Fast Timeout [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that have the specified fast-timeout configuration
setting.

``[-slow-timeout <integer>]`` - Slow Timeout [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that have the specified slow-timeout setting.

``[-mailbox-timeout <integer>]`` - Mailbox Timeout [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified mailbox-timeout setting.

``[-connect-timeout <integer>]`` - Connect Timeout [.advaced-
priv]#(privilege: advanced)#::
```

```
Selects the nodes that have the specified connect-timeout setting.

``[-operator-timeout <integer>]`` - Operator Timeout [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified operator-timeout setting.

``[-firmware-timeout <integer>]`` - Firmware Timeout [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified firmware-timeout setting.

``[-dumpcore-timeout <integer>]`` - Dumpcore Timeout [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified dump-core timeout setting.

``[-booting-timeout <integer>]`` - Booting Timeout [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified booting-timeout setting.

``[-transit-timer {true|false}]`` - Transit Timer Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified transit-timer setting.

``[-transit-timeout <integer>]`` - Transit Timeout [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified transit timeout.

``[-firmware-received <integer>]`` - Firmware Received [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified firmware-reception time.

``[-firmware-received-cycles <integer>]`` - Firmware Received in CPU
Cycles [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified firmware-reception time in CPU
Cycles.

``[-booting-received <integer>]`` - Booting Received [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified booting-reception time.

``[-transit-time <integer>]`` - Transit Event Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes whose last failover event occurred at the specified
time.

``[-pnormal <integer>]`` - Primary Normal IO Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified normal primary-mailbox I/O time.
```

```
``[-ptransition <integer>]`` - Primary Transition IO Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified transitional primary-mailbox I/O
time.

``[-bnormal <integer>]`` - Backup Normal IO Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified normal backup-mailbox I/O time.

``[-btransition <integer>]`` - Backup Transition IO Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified transitional backup-mailbox I/O
time.

``[-logs-unsynced <integer>]`` - Logs Unsynced Count [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified count of unsynchronized logs.

``[-auto-giveback {true|false}]`` - Auto Giveback Enabled::
Selects the nodes that have the specified auto-giveback setting.

``[-detection-time <integer>]`` - Takeover Detection Time (secs)::
Selects the nodes that have the specified detection-time setting.

``[-onfailure {true|false}]`` - Takeover on Failure Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified takeover-on-failure setting.

``[-onpanic {true|false}]`` - Takeover on Panic Enabled::
Selects the nodes that have the specified takeover-on-panic setting.

``[-onshort-uptime {true|false}]`` - Takeover on Short Uptime Enabled
[.advaced-priv]#(privilege: advanced)#::
Selects the storage-failover configurations that match this parameter
value.

``[-short-uptime <integer>]`` - Short Uptime (secs) [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified short-uptime value.

``[-attempts <integer>]`` - Number of Giveback Attempts [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified number of giveback attempts.

``[-attempts-time <integer>]`` - Giveback Attempts Period (minutes)
[.advaced-priv]#(privilege: advanced)#::
```

```
Selects the nodes that have the specified time setting for giveback
attempts.

``[-propagate {true|false}]`` - Propagate Status via Mailbox [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified propagate-status-via-mailbox
setting.

``[-read-interval <integer>]`` - Node Status Read Interval (secs)
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified read interval.

``[-write-interval <integer>]`` - Node Status Write Interval (secs)
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified write interval.

``[-onreboot {true|false}]`` - Takeover on Reboot Enabled::
Selects the nodes that have the specified takeover-on-reboot setting.

``[-delay-seconds <integer>]`` - Delay Before Auto Giveback (secs)::
Selects the nodes that have the specified delay (in seconds) for the auto
giveback.

``[-hwassist {true|false}]`` - Hardware Assist Enabled::
Selects the nodes that have the specified hwassist setting.

``[-hwassist-partner-ip <IP Address>]`` - Partner's Hwassist IP::
Selects the nodes that have the specified hwassist-partner-ip setting.

``[-hwassist-partner-port <integer>]`` - Partner's Hwassist Port::
Selects the nodes that have the specified hwassist-partner-port setting.

``[-hwassist-health-check-interval <integer>]`` - Hwassist Health Check
Interval (secs)::
Selects the nodes that have the specified hwassist health check interval,
in seconds.

``[-hwassist-retry-count <integer>]`` - Hwassist Retry Count::
Selects the nodes that have the specified hwassist retry count, in
seconds.

``[-hwassist-status <text>]`` - Hwassist Status::
Selects the nodes that have the specified hwassist-status setting.

``[-time-until-autogiveback <integer>]`` - Time Until Auto Giveback
(secs)::
Selects the nodes that have the specified time(in seconds) until auto
```

```
giveback.

``[-local-mailbox-disks <text>]`` - Local Mailbox Disks::
Selects the nodes that have the specified mailbox disks on the local node.

``[-partner-mailbox-disks <text>]`` - Partner Mailbox Disks::
Selects the nodes that have the specified mailbox disks on the partner
node.

``[-local-firmware-state <text>]`` - Local Firmware State [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified firmware state on the local
node.

``[-local-firmware-progress <integer>]`` - Local Firmware Progress Counter
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified firmware progress counter for
the local node.

``[-partner-firmware-state <text>]`` - Partner Firmware State [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified firmware state of the partner
node.

``[-partner-firmware-progress <integer>]`` - Partner Firmware Progress
Counter [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified firmware progress counter for
the partner node.

``[-local-missing-disks <text>]`` - Missing Disks on Local Node::
Selects the nodes that have the specified missing disks on the local node.

``[-partner-missing-disks <text>]`` - Missing Disks on Partner Node::
Selects the nodes that have the specified missing disks on the partner
node.

``[-reboot-timeout <integer>]`` - Reboot Timeout [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified reboot timeout.

``[-time-since-takeover <text>]`` - Time Since Takeover::
Selects the nodes that have been in takeover mode for the specified amount
of time.

``[-auto-giveback-after-panic {true|false}]`` - Auto Giveback After
Takeover On Panic::
Selects the nodes that have the specified auto-giveback-after-panic
```

```
setting. If ``_true_`` then an automatic giveback operation is invoked
when the failover node of an HA pair is repaired and rebooted. The
takeover node of the HA pair detects this and initiates a giveback
operation automatically.

``[-is-giveback-requested {true|false}]`` - Giveback Requested [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified is-giveback-requested setting.
If ``_true_`` , a deferred giveback request has been made by the local
node.

``[-auto-giveback-last-veto-check <integer>]`` - Auto Giveback Last Veto
Check [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified auto-giveback-last-veto-check
time. This setting indicates the time, in milliseconds, when the local
node made the most recent giveback veto check.

``[-is-auto-giveback-attempts-exceeded {true|false}]`` - Auto Giveback
Attempts Exceeded [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified is-auto-giveback-attempts-
exceeded setting. If ``_true_`` , the local node has exceeded the maximum
number of allowed auto giveback attempts.

``[-was-auto-giveback-done {true|false}]`` - Was Auto Giveback Done
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified was-auto-giveback-done setting.
If ``_true_`` , the last giveback was automatic (as opposed to a manual
giveback).

``[-is-cifs-auto-giveback-stopping {true|false}]`` - Is CIFS Auto Giveback
Stopping [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified is-cifs-auto-giveback-stopping
setting. If ``_true_`` , the local node has initiated CIFS termination as
part of an automatic giveback.

``[-aggregate-migration-timeout <integer>]`` - Aggregate Migration Timeout
(secs) [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified aggregate migration timeout.

``[-is-mirror-enabled {true|false}]`` - Is NVRAM Mirroring Enabled
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified is-mirror-enabled setting. If
``_true_`` , then NVRAM mirroring is enabled.

``[-is-mirror-consistency-required {true|false}]`` - Is Mirror Consistency
Required [.advaced-priv]#(privilege: advanced)#::
```

```
Selects the nodes that have the specified is-mirror-consistency-required
setting. If ``_true_`` , then NVRAM mirror consistency is required.

``[-is-memory-insufficient {true|false}]`` - Is Memory Insufficient To
Takeover [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified is-memory-insufficient setting.
If ``_true_`` , the local node does not have enough memory to perform a
takeover.

``[-memio-state <memio status>]`` - Current State of Memio Link [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified memio layer link current state.
Possible values are:
+
  * UNINIT - Uninitialized
  * CLOSED - Closed
  * HB_LISTEN - Listening for connect
  * SYN_SENT - Sent generation information
  * ESTABLISHED - Connection established



``[-is-degraded {true|false}]`` - Are Partner Mailbox Disks Not Known
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified is-degraded setting. If
``_true_`` , takeovers are deferred because partner mailbox disks are not
known.

``[-reserve-policy <reserve policy>]`` - FM Reservation Policy [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified disk reservation policy.
Possible values are:
+
  * RESERVE_NO_DISKS - no disk reservations made during takeover, nor are
disk reservations released during giveback
  * RESERVE_LOCK_DISKS_ONLY - only mailbox disks are released during
takeover and released during giveback
  * RESERVE_ONLY_AT_TAKEOVER - reservations are issued only at takeover
time. All disks are reserved. All reservations are released at giveback
  * RESERVE_ALWAYS_AFTER_TAKEOVER - reservations are issued at at
takeover. When disks are subsequently added, they are also reserved. All
disks are released at giveback



``[-reset-disks {true|false}]`` - Issue Disk Resets during Failover
[.advaced-priv]#(privilege: advanced)#::
```

```
Selects the nodes that have the specified reset-disks setting. If
``_true_`` , disks are reset during takeover/giveback.

``[-total-system-uptime <integer>]`` - Total System Uptime [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified total system uptime, in
milliseconds.

``[-current-time <integer>]`` - Current System Time [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that have the specified current time on the node.

``[-fm-takeover-state <FM Takeover/Giveback Transition>]`` - FM Takeover
State [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified takeover state. Possible values
are:
+
  * FT_NONE - Not in takeover
  * FT_TAKEOVER_STARTED - Local node has initiated takeover
  * FT_TAKEOVER_COMMITTED - Takeover has been committed
  * FT_TAKEOVER_DONE_OK - Local node successfully completed takeover
  * FT_TAKEOVER_DONE_FAILED - Takeover failed



``[-fm-giveback-state <FM Takeover/Giveback Transition>]`` - FM Giveback
State [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified giveback state. Possible values
are:
+
  * FT_NONE - Not in giveback
  * FT_GIVEBACK_READY - Partner node is ready for giveback
  * FT_GIVEBACK_STARTED - Local node has initiated giveback
  * FT_GIVEBACK_COMMITTED - Giveback has been committed
  * FT_GIVEBACK_DONE_OK - Giveback completed successfully



``[-takeover-reason <FM Takeover Reason>]`` - Reason why takeover
triggered [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that have the specified takeover reason. Possible values
are:
+
  * TAKEOVER_NONE - Not in takeover
  * TAKEOVER_IMMEDIATE - Operator initiated forced takeover
  * TAKEOVER_NDU - Takeover initiated as part of NDU
  * TAKEOVER_FORCED - Operator initiated forced takeover, possible data
```

```
loss
  * TAKEOVER_EARLY - Takeover occurred during the boot process
  * TAKEOVER_OPERATOR_EXP - Takeover occurred after the operator timeout
expired
  * TAKEOVER_POST_FAILED - Takeover occurred on POST failure
  * TAKEOVER_PANIC - Takeover on panic
  * TAKEOVER_SHORTUPTIME - Takeover after rapid toggling between up and
down states
  * TAKEOVER_SPARECORE_EXP - Takeover on panic timeout expiration
  * TAKEOVER_REBOOT_EXP - Takeover on reboot timer expiration
  * TAKEOVER_BOOTING_EXP - Takeover on booting timer expiration
  * TAKEOVER_FIRMWARE_EXP - Takeover on firmware timer expiration
  * TAKEOVER_NFO_SHUTDOWN - Takeover on negotiated failover shutdown
  * TAKEOVER_NFO_TIMER - Takeover on negotiated failover timer expiration
  * TAKEOVER_MDP - Takeover on multi-disk panic
  * TAKEOVER_REBOOT - Takeover on reboot
  * TAKEOVER_HALT - Takeover on halt
  * TAKEOVER_CLAM - CLAM-triggered takeover
  * TAKEOVER_HWASSIST - Hardware-assisted takeover
  * TAKEOVER_NORMAL - Operator initiated takeover



``[-ha-type {none|shared_storage|non_shared_storage}]`` - HA Type::
If this parameter is specified, the command selects the nodes that have
the specified HA-type setting. If the value is set to ``_shared_storage_``
, then the node is in a storage-failover configuration using the shared
storage. If it is set to ``_non_shared_storage_`` , then the node is in a
storage-failover configuration using the unshared storage. If it is set to
``_none_`` , then the node is not part of a storage-failover
configuration.

== Examples
```

```
The following example displays information about all storage-failover
configurations:
```

cluster1::> storage failover show
Takeover
Node Partner Possible State
-------- -------- -------- -----------------
node0 node1 true Connected to node1
node2 node3 true Connected to node3
node1 node0 true Connected to node0
node3 node2 true Connected to node2
4 entries were displayed.

```
[[ID5cbdf79a420e189f99304b417fe2bd4]]
= storage failover takeover
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Take over the storage of a node's partner

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover takeover` command initiates a takeover of the
partner node's storage.

== Parameters

{ ``-ofnode {<nodename>|local}`` - Node to Takeover::
This specifies the node that is taken over. It is shut down and its
partner takes over its storage.

| ``-bynode {<nodename>|local}`` - Node Initiating Takeover }::
This specifies the node that is to take over its partner's storage.

``[-option <takeover option>]`` - Takeover Option::
This optionally specifies the style of takeover operation. Possible values
include the following:
+
  * normal - Specifies a normal takeover operation; that is, the partner
is given the time to close its storage resources gracefully before the
takeover operation proceeds. This is the default value.
  * immediate - Specifies an immediate takeover. In an immediate takeover,
the takeover operation is initiated before the partner is given the time
to close its storage resources gracefully. The use of this option results
in an immediate takeover which does not do a clean shutdown. In case of
NDU this can result in a NDU failure.
+
NOTE: If this option is specified, negotiated takeover optimization is
bypassed even if the `-bypass-optimization` option is set to false.
+

+
NOTE: If this option is specified, migration of data LIFs from the partner
```

```
will be delayed even if the `-skip-lif-migration-before-takeover` option
is not specified. If possible, migrate the data LIFs to another node prior
to specifying this option.
+

  * allow-version-mismatch - If this value is specified, the takeover
operation is initiated even if the partner is running a version of
software that is incompatible with the version running on the node. In
this case, the partner is given the time to close its storage resources
gracefully before the takeover operation proceeds. However, the takeover
operation will not be allowed if the partner has higher WAFL or RAID label
versions. Use this value as part of a nondisruptive upgrade or downgrade
procedure.
  * force - If this value is specified, the takeover operation is
initiated even if the node detects an error that normally prevents a
takeover operation from occurring. This value is available only at the
advanced privilege level and higher.
+
NOTE: If this option is specified, negotiated takeover optimization is
bypassed even if the `-bypass-optimization` option is set to false.
+

+
NOTE: The use of this option can potentially result in data loss. If the
HA interconnect is detached or inactive, or the contents of the failover
partner's NVRAM cards are unsynchronized, takeover is normally disabled.
Using the `-force` option enables a node to take over its partner's
storage despite the unsynchronized NVRAM, which can contain client data
that can be lost upon storage takeover.
+




``[-bypass-optimization {true|false}]`` - Bypass Takeover Optimization
[.advaced-priv]#(privilege: advanced)#::
If this is an operator-initiated planned takeover, this parameter
specifies whether the takeover optimization is bypassed. This parameter
defaults to false.
+
NOTE: This parameter is ignored and negotiated takeover optimization
automatically bypassed if the `-immediate` option, the `-force` option, or
the `-allow-disk-inventory-mismatch` parameter is specified as part of the
same `storage failover takeover` command.
+
```

```
``[-allow-disk-inventory-mismatch {true|false}]`` - Disk inventory::
If this parameter is specified, the takeover operation is initiated even
if the local node cannot see the partner's filesystem disks.
+
NOTE: If this parameter is specified, negotiated takeover optimization is
bypassed even if the `-bypass-optimization` parameter is set to false.
+

+
NOTE: The use of this parameter can potentially result in client outage.
+



``[-skip-lif-migration-before-takeover <true>]`` - Skip Migrating LIFs
Away from Node Prior to Takeover::
This parameter specifies that LIF migration prior to takeover is skipped.
However if LIFs on this node are configured for failover, those LIFs may
still failover after the takeover has occurred. Without this parameter,
the command attempts to synchronously migrate data and cluster management
LIFs away from the node prior to its takeover. If the migration fails or
times out, the takeover is aborted.

``[-ignore-quorum-warnings <true>]`` - Skip Quorum Check Before Takeover::
If this parameter is specified, quorum checks will be skipped prior to the
takeover. The operation will continue even if there is a possible data
outage due to a quorum issue.

``[-override-vetoes <true>]`` - Override Vetoes::
If this is an operator-initiated planned takeover, this parameter
specifies whether the veto should be overridden. If this parameter is not
specified, its value is set to `false` .
+
NOTE: If this parameter is specified, negotiated takeover will override
any vetoes to continue with takeover.
+

+
NOTE: The use of this parameter might result in the takeover proceeding
even if the node detects issues that can potentially make the takeover
dangerous or disruptive.
+



``[-halt <true>]`` - Halt the Node That Is Taken Over::
This parameter specifies whether the node being taken over should be
halted. If the value is true, then the node being taken over is halted. If
```

```
the value is false, then the node being taken over is shutdown and might
be rebooted if AUTOBOOT is set to true. This parameter defaults to false.

``[-check-only <true>]`` - Only Run Pre-Takeover Checks::
This parameter initiates a verification check of a possible future planned
takeover. The check operation records any failures or issues that would
prevent a takeover. Use the xref:{relative_path}storage-failover-check-
takeover.html[storage failover check-takeover] command to view the result
of the check. This parameter defaults to false.
+
NOTE: This is a best effort operation. All checks passed does not
guarantee a successful planned takeover as failures may occur during
runtime.
+


== Examples

The following example causes a node named node0 to initiate a negotiated
optimized takeover of its partner's storage:
```

cluster1::> storage failover takeover -bynode node0

```
The following example causes a node named node0 to initiate an immediate
takeover of its partner's storage:
```

cluster1::> storage failover takeover -bynode node0 -option immediate

```
== Related Links

* xref:{relative_path}storage-failover-check-takeover.html[storage
failover check-takeover]




[[ID89f550c4db712054470a38e2099f5ec2]]
= storage failover hwassist show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display hardware-assisted storage failover status
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage failover hwassist show` command displays information about
the hardware-assisted storage failover status on each node. By default,
the command displays the following information:

* Node name.
* Partner node name.
* Whether hardware-assisted failover is enabled.
* IP address on which the local node receives hardware-assisted failover
alerts.
* Port on which the local node receives hardware-assisted failover alerts.
* Hardware-assisted failover status.
* If the monitor is inactive, the reason it is inactive.
* If the monitor is inactive, the corrective action to make it active.
* Status of keep-alive alerts on the local node.

Hardware-assisted failover establishes a notification channel from each
respective node's service processor to the other (HA partner) node. If a
node becomes unresponsive, its service processor notifies the HA partner
of this condition, accelerating storage failover. By default, hwassist is
enabled and configured automatically to use each node's node-mgmt LIF. To
modify or show the hardware-assisted storage failover configuration, use
the xref:{relative_path}storage-failover-modify.html[storage failover
modify] and xref:{relative_path}storage-failover-show.html[storage
failover show] commands.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays information only about
hwassist on the specified nodes.
```

```
``[-partner-name {<nodename>|local}]`` - Name of the Partner Node::
If you specify this parameter, the command displays information only about
hwassist on nodes with the specified partner node.

``[-enabled {true|false}]`` - Local Hardware Assist Enabled::
If you specify this parameter, the command displays information only about
hwassist on nodes with the specified enabled state.

``[-local-status <text>]`` - Local Node's Hwassist Status::
If you specify this parameter, the command displays information only about
hwassist on nodes with the specified local status (``_active_`` or
``_inactive_`` ).

``[-local-ip <text>]`` - IP Address on Which Local Node is Listening::
If you specify this parameter, the command displays information only about
hwassist on nodes with the specified local IP address.

``[-local-port <integer>]`` - Port on Which Local Node is Listening::
If you specify this parameter, the command displays information only about
hwassist on nodes with the specified local UDP port.

``[-local-inactive <text>]`` - Local Node's Hwassist Inactive Status
Reason::
If you specify this parameter, the command displays information only about
hwassist on nodes with the specified inactive reason.

``[-local-action <text>]`` - Corrective Action on Local Node::
If you specify this parameter, the command displays information only about
hwassist on nodes with the specified corrective action.

== Examples

The following example displays the hardware-assisted failover information
for node cluster1-01 and its HA partner node cluster1-02:
```

cluster1::> storage failover hwassist show
Node
---
cluster1-01
Partner: cluster1-02
Hwassist Enabled: true
Hwassist IP: 10.225.248.19
Hwassist Port: 4444
Monitor Status: active
Inactive Reason: -
Corrective Action: -
Keep-Alive Status: healthy
cluster1-02

Partner: cluster1-01
Hwassist Enabled: true
Hwassist IP: 10.225.248.21
Hwassist Port: 4444
Monitor Status: active
Inactive Reason: -
Corrective Action: -
Keep-Alive Status: healthy

```
== Related Links

* xref:{relative_path}storage-failover-modify.html[storage failover
modify]
* xref:{relative_path}storage-failover-show.html[storage failover show]




[[IDf703dae59f5886c19d8638dd559858b8]]
= storage failover hwassist test
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Test the hwassist functionality

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The ` storage failover hwassist test` command tests the Hardware Assist
h/w connectivity between the two nodes in a HA pair. The test result can
be one of the following.

* Hardware Assist is not initialized.
* HW assist is not supported.
* Partner is throttling alerts.
* Resource is busy.
* Hardware Assist h/w returned an error.
* No response from partner.Timed out.
* Unexpected abort.
* Partner has taken over.
* Interconnect is down between nodes.
* Partner is not booted up yet.
```

```
== Parameters

``-node {<nodename>|local}`` - Node::
This specifies the node from which a test alert is initiated.

== Examples

The following command issues a test alert from the node cluster1-01:
```

cluster1::> storage failover hwassist test -node cluster1-01
Info: Operation successful.

```
[[ID11a16d0239c359a0eee540ad8c89aeff]]
= storage failover hwassist stats clear
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear the hwassist statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The ` storage failover hwassist stats clear` command clears the statistics
information maintained by Hardware Assist functionality.

== Parameters

``-node {<nodename>|local}`` - Node::
This specifies the node on which the statistics are to be cleared.

== Examples

The following example clears the hwassist statistics on the node cluster1-
01:
```

cluster1::> storage failover hwassist stats clear -node cluster1-01

```
[[ID d89e8585a2a8718eb6c07e7ad40716f1]]
= storage failover hwassist stats show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display hwassist statistics

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The` storage failover hwassist stats show ` command displays statistics
about the hardware assist alerts processed by a node. The command displays
the following information for each alert:

* Locally enabled.
* Partner Inactive Reason.
* Alert type.
* Event that triggered the alert.
* The number of times the alert has been received.
* Whether takeover was possible on receiving the alert.
* The last time at which the alert was received.



== 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>|local}]`` - Node::
Selects the hwassist statistics for the specified node.

``[-enabled {true|false}]`` - Local Hardware Assist Enabled::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified enabled state.
```

```
``[-local-inactive <text>]`` - Local Node's Hwassist Inactive Status
Reason::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified inactive reason.

``[-power-loss <text>]`` - Power Loss Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified power-loss event.

``[-watchdog-reset <text>]`` - Watchdog Reset Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified watchdog-reset event.

``[-power-off-via-rlm <text>]`` - Power Off by Using RLM::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified power off event via Remote
LAN Manager.

``[-power-cycle-via-rlm <text>]`` - Power Cycle by Using RLM::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified power cycle event via
Remote LAN Manager.

``[-reset-via-rlm <text>]`` - System Reset by Using RLM::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified reset event via Remote LAN
Manager.

``[-power-off-via-sp <text>]`` - Power Off by Using SP::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified power off event via
Service Processor.

``[-power-cycle-via-sp <text>]`` - Power Cycle by Using SP::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified power cycle event via
Service Processor.

``[-reset-via-sp <text>]`` - System Reset by Using SP::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified reset event via Service
Processor.

``[-dimm-uecc-error <text>]`` - DIMM UECC Fatal Error Detected by Storage
OS::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified DIMM error event.
```

```
``[-post-error <text>]`` - POST Error Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified post error event.

``[-abnormal-reboot <text>]`` - Abnormal Reboot Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified abnormal reboot event.

``[-loss-of-heartbeat <text>]`` - Loss of Heartbeat Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified loss of heartbeat event.

``[-keep-alive <text>]`` - Periodic Message Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified keep alive event.

``[-test <text>]`` - Test Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified test event.

``[-bad-nvram-id <text>]`` - Bad NVRAM ID Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified bad NVRAM ID event.

``[-ss-mismatch <text>]`` - Shared Secret Key Mismatch Alert::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified shared secret mismatch
event.

``[-unknown-alerts <text>]`` - Unknown Alerts::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified unknown alerts event.

``[-times-throttled <text>]`` - Number of Throttled Alerts::
If you specify this parameter, the command displays information only about
hwassist statistics on nodes with the specified times throttled event.

== Examples

The following example displays the hwassist statistics for the node ha1:
```

cluster1::> storage failover hwassist stats show -node ha1
Node: ha1
Local Enabled: true
Partner Inactive Reason: -
Alert Type Alert Event Count Takeover Last Received

---------- ---------- ------ -------- ----------
system_down power_loss 0 Yes ---
system_down l2_watchdog_reset 0 Yes ---
system_down power_off_via_rlm 0 Yes ---
system_down power_cycle_via_rlm 0 Yes ---
system_down reset_via_rlm 0 Yes ---
system_down power_off_via_sp 0 Yes ---
system_down power_cycle_via_sp 0 Yes ---
system_down reset_via_sp 0 Yes ---
system_down dimm_uecc_error 0 Yes ---
system_down post_error 0 No ---
system_down abnormal_reboot 0 No ---
system_down loss_of_heartbeat 0 No ---
keep_alive periodic_message 10 No Wed Mar 9 13:41:28 EST 2016
test test 0 No ---
ID_mismatch --- 0 --- ---
Key_mismatch --- 0 --- ---
Unknown --- 0 --- ---
alerts_throttled 0 --- ---

```
The following example displays the hwassist statistics for the node ha1
where hardware assist hardware is not supported.
```

cluster1::> storage failover hwassist stats show -node ha1
Node: ha1
Local Enabled: false
Partner Inactive Reason: HW assist is not supported on partner.
Alert Type Alert Event Count Takeover Last Received
---------- ---------- ------ -------- ----------
-

```
[[ID805cd86814e6e5e34c5ffbf3d2d399f5]]
= storage failover internal-options show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the internal options for storage failover

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage failover internal-options show` command displays the
following information about the storage failover configuration:
```

```
* Node name
* Whether automatic giveback is enabled
* Whether partner checking is enabled
* Takeover detection time, in seconds
* Whether takeover on failover is enabled
* Whether takeover on panic is enabled
* Whether takeover on reboot is enabled
* Whether hardware-assisted takeover is enabled
* IP address on which the partner node listens to the hardware-assisted
takeover alerts
* Port on which the partner node listens to the hardware-assisted takeover
alerts
* Whether takeover on short uptime is enabled (detailed view only)
* Short uptime interval, in seconds (detailed view only)
* Number of giveback attempts (detailed view only)
* Giveback attempt interval, in minutes (detailed view only)
* Whether status is propagated through SFO mailboxes (detailed view only)
* Status read interval, in seconds (detailed view only)
* Status write interval, in seconds (detailed view only)
* Hardware-assisted takeover retry count (detailed view only)
* Hardware-assisted takeover heartbeat period (detailed view only)
* Whether operator-initiated planned takeover is optimized



== 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.

| ``[-more  ]`` [.advaced-priv]#(privilege: advanced)#::
This parameter displays the following additional information: :
+
  * Node name
  * Whether takeover on short uptime is enabled
  * Short uptime interval, in seconds
  * Number of giveback attempts
  * Giveback attempt interval, in minutes
  * Whether status is propagated through SFO mailboxes
  * Status read interval, in seconds
  * Status write interval, in seconds
  * Hardware-assisted takeover retry count
  * Hardware-assisted takeover heartbeat period
```

```
| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects configuration information for the specified node.

``[-auto-giveback {true|false}]`` - Auto Giveback Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
automatic giveback setting.

``[-check-partner {true|false}]`` - Check Partner Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
partner-checking setting.

``[-detection-time <integer>]`` - Takeover Detection Time (secs)
[.advaced-priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
takeover detection time setting.

``[-onfailure {true|false}]`` - Takeover on Failure Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
takeover-on-failure setting.

``[-onpanic {true|false}]`` - Takeover on Panic Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
takeover-on-panic setting.

``[-onshort-uptime {true|false}]`` - Takeover on Short Uptime Enabled
[.advaced-priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
takeover-on-short-uptime setting.

``[-short-uptime <integer>]`` - Short Uptime (secs) [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
takeover-on-short-uptime time setting.

``[-attempts <integer>]`` - Number of Giveback Attempts [.advaced-
priv]#(privilege: advanced)#::
```

```
Selects configuration information for nodes that have the specified number
of giveback attempts setting.

``[-attempts-time <integer>]`` - Giveback Attempts Minutes [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
giveback attempt time setting.

``[-propagate {true|false}]`` - Propagate Status via Mailbox [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
setting for propagation of status through Storage Failover mailboxes.

``[-read-interval <integer>]`` - Node Status Read Interval (secs)
[.advaced-priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified status
read interval setting.

``[-write-interval <integer>]`` - Node Status Write Interval (secs)
[.advaced-priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified status
write interval setting.

``[-onreboot {true|false}]`` - Takeover on Reboot Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
takeover-on-reboot setting.

``[-delay-seconds <integer>]`` - Delay Before Auto Giveback (secs)
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the node or nodes that have the specified delay for auto giveback.

``[-hwassist {true|false}]`` - Hwassist Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
hardware-assisted takeover setting.

``[-hwassist-partner-ip <text>]`` - Partner's Hwassist IP [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
partner IP setting for hardware-assisted takeovers.

``[-hwassist-partner-port <integer>]`` - Partner's Hwassist Port
[.advaced-priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
partner port setting for hardware-assisted takeovers.
```

```
``[-hwassist-health-check-interval <integer>]`` - Hwassist Health Check
Interval (secs) [.advaced-priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified health
check interval setting for hardware-assisted takeovers

``[-hwassist-retry-count <integer>]`` - Hwassist Retry Count [.advaced-
priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified retry
count (in seconds) for hardware-assisted takeovers.

``[-mode {ha|non_ha}]`` - HA Mode [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the node or nodes that have the specified HA mode.

``[-bypass-takeover-optimization {true|false}]`` - Bypass Takeover
Optimization Enabled [.advaced-priv]#(privilege: advanced)#::
Selects configuration information for nodes that have the specified
setting for bypass takeover optimization ( ``_ true_`` means that
optimized operator-initiated planned takeover is bypassed, ``_false_``
means that it is enabled). Operator-initiated planned takeover is
optimized when SFO aggregates are relocated serially to the partner prior
to takeover. This reduces client outage.

== Examples

The following example displays detailed information about the internal
options for storage failover on a node named node2:
```

cluster1::*> storage failover internal-options show -node node2
Node: node2
Auto Giveback Enabled: false
Check Partner Enabled: true
Takeover Detection Time (secs): 15
Takeover On Failure Enabled: true
Takeover On Panic Enabled: false
Takeover On Short Uptime Enabled: true
Short Uptime (secs): -
Number of Giveback Attempts: 3
Giveback Attempts Minutes: 10
Propagate Status Via Mailbox: true
Node Status Read Interval (secs): 5
Node Status Write Interval (secs): 5
Failover the Storage when Cluster Ports Are Down: -
Failover Interval when Cluster Ports Are Down (secs): -
Takeover on Reboot Enabled: true
Delay Before Auto Giveback (secs): 300
Hardware Assist Enabled: true
Partner’s Hw-assist IP:

Partner's Hw-assist Port: 4444
Hw-assist Health Check Interval (secs): 180
Hw-assist Retry count: 2
HA mode: ha
Bypass Takeover Optimization Enabled: true

```
[[ID3ea387bc6ba382578464e2fdf3c181b6]]
= storage failover mailbox-disk show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about storage failover mailbox disks

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage failover mailbox-disk show` command lists the mailbox disks
that are used by storage failover. The command displays the following
information:

* Node name
* Whether the mailbox disk is owned by the local node or by its partner
* Disk name
* Disk universal unique identifier (UUID)

This command is available only at the advanced privilege level and higher.

== Parameters

``[-fields <fieldname>,...]``::
If -fields <fieldname>,... is used, the command displays only the
specified fields.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is used, the command displays detailed information about
all entries.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects the mailbox disks that are associated with the specified node.

``[-location {local|partner|tertiary}]`` - Mailbox Owner [.advaced-
```

```
priv]#(privilege: advanced)#::
Selects the mailbox disks that have the specified relationship to the
node.

``[-diskindex <integer>]`` - Mailbox Disk Index [.advaced-
priv]#(privilege: advanced)#::
Selects the mailbox disk that has the specified index number.

``[-diskname <text>]`` - Mailbox Disk Name [.advaced-priv]#(privilege:
advanced)#::
Selects the mailbox disks that match the specified disk name.

``[-diskuuid <text>]`` - Mailbox Disk UUID [.advaced-priv]#(privilege:
advanced)#::
Selects the mailbox disks that match the specified UUID.

``[-physical-location {local|partner|mediator}]`` - Mailbox Disk Physical
Location [.advaced-priv]#(privilege: advanced)#::
Selects the mailbox disks that match the specified physical location.

``[-location-id <nvramid>]`` - System ID of the Node where the Disk is
Attached [.advaced-priv]#(privilege: advanced)#::
Selects the mailbox disks that match the specified location-id.

``[-location-name <text>]`` - Mailbox Disk Location [.advaced-
priv]#(privilege: advanced)#::
Selects the mailbox disks that match the specified location-name.

== Examples

The following example displays information about the mailbox disks on a
node named node1:
```

cluster1::*> storage failover mailbox-disk show -node node1
Node Location Index Disk Name Physical Location Disk UUID
------- --------- ----- ------------ ----------------- -------------------
node1
local 0 1.0.4 local 20000000:8777E9D6:[…]
local 1 1.0.6 partner 20000000:8777E9DE:[…]
partner 0 1.0.1 local 20000000:877BA634:[…]
partner 1 1.0.2 partner 20000000:8777C1F2:[…]

```
[[ID5e77e9f1fafc092091e636f7b20ca91e]]
= storage failover progress-table show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status information about storage failover operations

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage failover progress-table show` displays status information
about storage-failover operations. This information is organized in a
resource table. The command displays the following information:

* Node name
* Resource-entry index number
* Resource-entry name
* Resource-entry state
* Resource-entry failure code
* Resource-entry time delta

This command is available only at the advanced privilege level and higher.

== Parameters

``[-fields <fieldname>,...]``::
If -fields <fieldname>, ... is used, the command will only displays only
the specified fields.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is used, the command displays detailed information about
all entries.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects the status information for the specified node.

``[-index <integer>]`` - Resource Table Index [.advaced-priv]#(privilege:
advanced)#::
Selects the status information for the specified index number.

``[-entryname <text>]`` - Resource Table Entry Name [.advaced-
priv]#(privilege: advanced)#::
Selects the status information for the specified entry name.

``[-state <text>]`` - Resource Table Entry State [.advaced-
priv]#(privilege: advanced)#::
```

```
Selects the status information for the specified state. Possible values
include UP, START_RUNNING, START_DONE, START_FAILED, STOP_RUNNING,
STOP_FAILED, TAKEOVER_BARRIER, and ONLY_WHEN_INITD.

``[-failurecode <text>]`` - Entry Failure Code [.advaced-priv]#(privilege:
advanced)#::
Selects the status information for the specified failure code. Possible
values include OK, FAIL, FAIL_ALWAYS, HANG, PANIC, and VETO.

``[-timedelta <integer>]`` - Entry Time Delta [.advaced-priv]#(privilege:
advanced)#::
Selects the status information for the specified time delta.

== Examples

The following example displays the entire storage-failover resource table:
```

cluster1::*> storage failover progress-table show
Node Entry Name State Time Delta
------ ---------- ------- ----------
node0
Pre-rsrctbl: fmdisk_resumePartnerDi start_done 6
Pre-rsrctbl: coredump_get_busy_spar start_done 107
Pre-rsrctbl: raid_preread_labels_be start_done 1
Pre-rsrctbl: fmdisk_reserve_all start_done 84
rsrctbl: fmrsrc_giveback_done start_done 0
rsrctbl: fmic start_done 0
rsrctbl: fmdisk_reserve start_done 171
rsrctbl: fm_partnerSlowTimeout start_done 1
rsrctbl: fmdisk_inventory start_done 0
rsrctbl: fmfsm_reserve start_done 0
Press <space> to page down, <return> for next line, or 'q' to quit…
Node Entry Name State Time Delta
------ ---------- ---------- ----------
node0
rsrctbl: rdb-ha start_done 36
rsrctbl: giveback_cleanup_wait start_done 0
rsrctbl: priority_ha start_done 0
rsrctbl: raid start_done 113
rsrctbl: raid_disaster_early start_done 0
rsrctbl: wafl_nvram_replay start_done 0
rsrctbl: takeover_test_1 start_done 0

```
:leveloffset: -1



= storage firmware commands

:leveloffset: +1
```

```
[[ID2c6bfb68f6228931657f6435fecfaea3]]
= storage firmware download
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Download disk, ACP processor and shelf firmware

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage firmware download` command downloads the ACP processor, disk,
and shelf module firmware to a specified node. This command can also be
used to download the disk qualification package (DQP).

Use the xref:{relative_path}storage-disk-firmware-update.html[storage disk
firmware update] command to install downloaded disk firmware.

Use the xref:{relative_path}storage-shelf-firmware-update.html[storage
shelf firmware update] command to install downloaded shelf module
firmware.

Use the xref:{relative_path}system-node-run.html[system node run]` local
storage download acp` command to install downloaded ACP processor
firmware.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the node to which the firmware is to be downloaded.

``-package-url <text>`` - Package URL [.advaced-priv]#(privilege:
advanced)#::
This specifies the path to the firmware package.
+
The packaged ACP processor, disk, and shelf firmware files need to have
".AFW", ".LOD", and ".SFW" file extensions, respectively.
+
The following URL protocols are supported: ftp, http, tftp and file. The
file URL scheme can be used to specify the location of the package to be
```

```
fetched from an external device connected to the storage controller.
Currently, only USB mass storage devices are supported. The USB device is
specified as ``_\file://usb0/<filename>_`` . The package must be present
in the root directory of the USB mass storage device.

== Examples

The following example downloads a disk firmware package with the path
``_\ftp://example.com/fw/disk-fw-1.2.LOD.zip_`` to a node named
``_node1_`` :
```

cluster1::> storage firmware download -node node1 -package-url ftp://example.com/fw/disk-fw-1.2.LOD.zip

```
== Related Links

* xref:{relative_path}storage-disk-firmware-update.html[storage disk
firmware update]
* xref:{relative_path}storage-shelf-firmware-update.html[storage shelf
firmware update]
* xref:{relative_path}system-node-run.html[system node run]




[[IDa9a6cb9a2bb804b1be4988a9a85ea0e2]]
= storage firmware acp delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an ACP firmware file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage firmware acp delete` command deletes the specified ACP
processor firmware file from all nodes that are currently part of the
cluster.

== Parameters

``-filename <text>`` - Firmware Filename [.advaced-priv]#(privilege:
advanced)#::
Specifies the firmware file to delete.

== Examples

The following example deletes the ACP processor firmware file with the
name ACP-IOM3.0150.AFW.FVF on each node:
```

cluster1::*> storage firmware acp delete -filename ACP-IOM3.0150.AFW.FVF

```
[[ID072d4a92fdf192d2965490681bef04a5]]
= storage firmware acp rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename an ACP firmware file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage firmware acp rename` command renames the specified ACP
processor firmware file on each node.

== Parameters

``-oldname <text>`` - Old Filename [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the firmware file to rename.

``-newname <text>`` - New Filename [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the new name of the firmware file.

== Examples

The following example renames the ACP processor firmware file with the
name ACP-IOM3.0150.AFW.FVF to ACP-IOM3.AFW.FVF on each node:
```

cluster1::*> storage firmware acp rename -oldname ACP-IOM3.0150.AFW.FVF -newname ACP-IOM3.AFW.FVF

```
[[ID72f94a8f00c443c24a7a633dd6582706]]
= storage firmware acp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of ACP firmware files on the given node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage firmware acp show` command displays the ACP processor
firmware files present on each node.

== 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>|local}]`` - Node::
Selects the files that match the specified node name.

``[-filename <text>]`` - Storage Firmware File::
Selects the files that match the specified filename.

== Examples

The following example displays the ACP processor firmware files on each
node:
```

cluster1::> storage firmware acp show
Node: Node1
ACP Firmware Files

ACP-IOM3.0150.AFW.FVF
ACP-IOM3.AFW

ACP-IOM6.0210.AFW
ACP-IOM6.0210.AFW.FVF
Node: Node2
ACP Firmware Files

ACP-IOM3.0150.AFW.FVF
ACP-IOM3.AFW
ACP-IOM6.0210.AFW
ACP-IOM6.0210.AFW.FVF
8 entries were displayed.

```
[[IDb51bb4fe330784fd8b291236c55fbd43]]
= storage firmware disk delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a disk firmware file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage firmware disk delete` command deletes the specified disk
firmware file on each node.

== Parameters

``-filename <text>`` - Storage Firmware Filename [.advaced-
priv]#(privilege: advanced)#::
Specifies the firmware file to delete.

== Examples

The following example deletes the disk firmware file with the name
X262_SMOOST25SSX.NA06.LOD on each node:
```

cluster1::*> storage firmware disk delete -filename X262_SMOOST25SSX.NA06.LOD

```
[[IDeebcb02e0cc1d5370935545c54c87531]]
= storage firmware disk rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a disk firmware file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage firmware disk rename` command renames the specified disk
firmware file on each node.

== Parameters

``-oldname <text>`` - Old Filename [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the firmware file to rename.

``-newname <text>`` - New Filename [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the new name of the firmware file.

== Examples

The following example renames the disk firmware file with the name
X262_SMOOST25SSX.NA06.LOD to X262_SMOOST25SSX.LOD on each node:
```

cluster1::*> storage firmware disk rename -oldname X262_SMOOST25SSX.NA06.LOD -newname X262_SMOOST25SSX.LOD

```
[[IDd2f686f3af26a1dd0bc66a015da6d4f1]]
= storage firmware disk show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of disk firmware files on the given node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage firmware disk show` command displays the disk firmware files
present on each node.

== 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>|local}]`` - Node::
Selects the files that match the specified node name.

``[-filename <text>]`` - Storage Firmware File::
Selects the files that match the specified filename.

== Examples

The following example displays the disk firmware files on each node:
```

cluster1::> storage firmware disk show
Node: Node1
Disk Firmware Files

X262_SMOOST25SSX.NA06.LOD
X262_SMOOST25SSX.NA06.LOD.FVF
X267_SMOOST50SSX.NA06.LOD

```
X267_SMOOST50SSX.NA06.LOD.FVF
Node: Node2
Disk Firmware Files
----------------------------------------
X262_SMOOST25SSX.NA06.LOD
X262_SMOOST25SSX.NA06.LOD.FVF
X267_SMOOST50SSX.NA06.LOD
X267_SMOOST50SSX.NA06.LOD.FVF
8 entries were displayed.
```

```
[[IDefa7ff15b440a76a2eaac49cdf6ac245]]
= storage firmware hba show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of HBA firmware files on the given node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage firmware hba show` command displays the HBA firmware files
present on each node. For MCC configs, run the command on each DR group.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>, ...` parameter, the command
output also includes the specified field or fields. You can use '-fields
?' to display the fields to specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Selects the files that match the specified node name.

``[-filename <text>]`` - Storage Firmware File::
Selects the files that match the specified filename.

== Examples

The following example displays the HBA firmware files on each node:
```

cluster1::> storage firmware hba show
Node: Node1
HBA Firmware Files

PM8072.03080900.HFW
SAS3008.13000000.HFW
SAS3616W.12030003.HFW

PM8072.03080900.HFW
Node: Node2
HBA Firmware Files

PM8072.03080900.HFW
SAS3008.13000000.HFW
SAS3616W.12030003.HFW
PM8072.03080900.HFW
8 entries were displayed.

```
[[ID6955d8131ca203fb249aad040fd8c74c]]
= storage firmware shelf delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a shelf firmware file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage firmware shelf delete` command deletes the specified shelf
firmware file from all nodes that are currently part of the cluster.

== Parameters

``-filename <text>`` - Storage Firmware Filename [.advaced-
priv]#(privilege: advanced)#::
Specifies the firmware file to delete.

== Examples

The following example deletes the shelf firmware file with the name
IOM12.0210.SFW on each node:
```

cluster1::*> storage firmware shelf delete -filename IOM12.0210.SFW

```
[[ID8e65eabaa14d133a34b11c78f61dfaf3]]
= storage firmware shelf rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a shelf firmware file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage firmware shelf rename` command renames the specified shelf
firmware file on each node.

== Parameters

``-oldname <text>`` - Old Filename [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the firmware file to rename.

``-newname <text>`` - New Filename [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the new name of the firmware file.

== Examples

The following example renames the shelf firmware file with the name
IOM12.0210.SFW to IOM12.000.SFW on each node:
```

cluster1::*> storage firmware shelf rename -oldname IOM12.0210.SFW -newname IOM12.000.SFW

```
[[IDbc4e1b1d07018fd5d6af95c9ddd9752b]]
= storage firmware shelf show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of shelf firmware files on the given node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage firmware shelf show` command displays the shelf firmware
files present on each node.

== 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>|local}]`` - Node::
Selects the files that match the specified node name.

``[-filename <text>]`` - Storage Firmware File::
Selects the files that match the specified filename.

== Examples

The following example displays the shelf firmware files on each node:
```

cluster1::> storage firmware shelf show
Node: Node1
Shelf Firmware Files

AT-FCX.3800.SFW
AT-FCX.3800.SFW.FVF
ESH4.1400.SFW

ESH4.1400.SFW.FVF
Node: Node2
Shelf Firmware Files

AT-FCX.3800.SFW
AT-FCX.3800.SFW.FVF
ESH4.1400.SFW
ESH4.1400.SFW.FVF
8 entries were displayed.

```
:leveloffset: -1


= storage iscsi-initiator commands

:leveloffset: +1


[[ID5f687a36d4514f5862506e79d255800e]]
= storage iscsi-initiator add-target
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add an iSCSI target

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage iscsi-initiator add-target` command adds an iSCSI target to a
node's list of targets. This command is only supported on high-
availability shared-nothing virtualized platforms.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the ONTAP node to which the iSCSI target will be
added.

``-label <text>`` - User Defined Identifier [.advaced-priv]#(privilege:
advanced)#::
Specifies a label for the target to be added.
```

```
``-target-type
{external|mailbox|partner|partner2|dr_auxiliary|dr_partner}`` - Target
Type [.advaced-priv]#(privilege: advanced)#::
Specifies the type of the target. It is used by the node to determine how
to use the LUNs. There are five target types:
+
  * partner - The partner target should belong to the node's HA partner.
This allows the node to access its partner's disks.
  * mailbox - A mailbox target's LUNs are used exclusively as HA
mailboxes.
  * external - External targets' LUNs can be used by the node but do not
play a role in HA.
  * dr_auxiliary - The DR auxiliary target for MetroCluster over IP. Not a
valid target type for the add-target command.
  * dr_partner - The DR partner target for MetroCluster over IP. Not a
valid target type for the add-target command.



``-target-portal <text>`` - Target Portal [.advaced-priv]#(privilege:
advanced)#::
Specifies the target's IP address and listening TCP port. The port is not
required if it is the default iSCSI port (3260). Examples of correct
target portals are ``_10.0.0.2_`` and ``_10.0.0.2:860_`` .

``-target-name <text>`` - iSCSI Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the iSCSI target name such as an IQN (iSCSI qualified name).

``[-status-admin {down|up}]`` - Administrative Status (default: up)
[.advaced-priv]#(privilege: advanced)#::
Use to specify whether the initial administrative status of the connection
is up or down. The default setting is ``_up_`` .

== Examples

The following example adds and connects to an iSCSI target from the
specified node.
```

cluster1::*> storage iscsi-initiator add-target -node node1
-label target1 -target-type external
-target-portal 10.0.0.2:860
-target-name iqn.2012-06.com.bsdctl:target0

```
[[IDdec802d5b937c7d1d5f98d3e940bdfc8]]
= storage iscsi-initiator connect
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Connect to an iSCSI target

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage iscsi-initiator connect` command connects a node to the
specified target. This command is only supported on high-availability
shared-nothing virtualized platforms.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the ONTAP node to which the iSCSI target will be
connected.

``[-target-type
{external|mailbox|partner|partner2|dr_auxiliary|dr_partner}]`` - Target
Type [.advaced-priv]#(privilege: advanced)#::
Selects targets with the specified target type.

``-label <text>`` - User Defined Identifier [.advaced-priv]#(privilege:
advanced)#::
Specifies the label of the target to connect to.

== Examples

The following example adds and connects to an iSCSI target from the
specified node.
```

cluster1::*> storage iscsi-initiator connect -node node1
-label target1

```
[[ID3034209a867c990ffd1e1fd839b303d5]]
= storage iscsi-initiator disconnect
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disconnect from an iSCSI target

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage iscsi-initiator disconnect` command disconnects a node from
the specified target. This command is only supported on high-availability
shared-nothing virtualized platforms.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the ONTAP node from which the iSCSI target will be
disconnected.

``[-target-type
{external|mailbox|partner|partner2|dr_auxiliary|dr_partner}]`` - Target
Type [.advaced-priv]#(privilege: advanced)#::
Selects targets with the specified target type.

``-label <text>`` - User Defined Identifier [.advaced-priv]#(privilege:
advanced)#::
Specifies the label of the target to disconnect from.

== Examples

The following example adds and connects to an iSCSI target from the
specified node.
```

cluster1::*> storage iscsi-initiator disconnect -node node1
-label target1

```
[[ID8dd04727349985ca0e936cf9aeb512cb]]
= storage iscsi-initiator remove-target
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove an iSCSI target

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage iscsi-initiator remove-target` command removes an iSCSI
target from a node's list of targets. This command is only supported on
high-availability shared-nothing virtualized platforms.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the ONTAP node from which the iSCSI target will be
removed.

``[-target-type
{external|mailbox|partner|partner2|dr_auxiliary|dr_partner}]`` - Target
Type [.advaced-priv]#(privilege: advanced)#::
Selects targets with the specified target type.

``-label <text>`` - User Defined Identifier [.advaced-priv]#(privilege:
advanced)#::
Specifies the label of the target to be removed.

== Examples

The following example adds and connects to an iSCSI target from the
specified ``_node_`` .
```

cluster1::*> storage iscsi-initiator remove-target -node node1
-label target1

```
[[ID4625c0e3ce42a1cd2d4ae374de059fbf]]
= storage iscsi-initiator show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the iSCSI targets

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage iscsi-initiator show` displays the list of iSCSI targets
configured for each ONTAP node in the cluster. This command is only
supported on high-availability shared-nothing virtualized platforms.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Represents the name of the ONTAP node for which information is to be
displayed. If this parameter is not specified, the command displays
information about all nodes in the cluster.

``[-target-type
{external|mailbox|partner|partner2|dr_auxiliary|dr_partner}]`` - Target
Type [.advaced-priv]#(privilege: advanced)#::
Selects targets with the specified target type.

``[-label <text>]`` - User Defined Identifier [.advaced-priv]#(privilege:
advanced)#::
Selects targets with the specified label.

``[-target-portal <text>]`` - Target Portal [.advaced-priv]#(privilege:
advanced)#::
Selects targets with the specified portal.
```

```
``[-target-name <text>]`` - iSCSI Name [.advaced-priv]#(privilege:
advanced)#::
Selects targets with the specified target name.

``[-status-admin {down|up}]`` - Administrative Status [.advaced-
priv]#(privilege: advanced)#::
Selects targets with the specified administrative status.

``[-status-oper {down|up}]`` - Operational Status [.advaced-
priv]#(privilege: advanced)#::
Selects targets with the specified operational status.

``[-failure-reason <text>]`` - Failure Reason [.advaced-priv]#(privilege:
advanced)#::
Selects targets with the specified failure reason.

== Examples

The following example displays the list of iSCSI targets for each node in
the cluster.
```

cluster1::*> storage iscsi-initiator show
Status
Node Type Label Target Portal Target Name Admin/Op
---- ---- -------- ---------- ---------- ---------- --------
node1
mailbox
mediator 10.235.14.141 iqn.2012-05.local:mailbox.group.1 up/up
partner
partner 10.63.7.205:65200 iqn.2012-06.com.bsdctl:target0 up/up
node2
mailbox
mediator 10.235.14.141 iqn.2012-05.local:mailbox.group.1 up/up
partner
partner 10.63.7.201:65200 iqn.2012-06.com.bsdctl:target0 up/up
4 entries were displayed.

```
:leveloffset: -1


= storage path commands

:leveloffset: +1


[[IDe841c992de742aedf074c97176220e52]]
= storage path quiesce
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Quiesce I/O on a path to array

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage path quiesce` command quiesces I/O on one path to a LUN. It
also quiesces the given entire path immediately or can monitor the given
path for error threshold before quiesce. After the I/O has been quiesced,
no new I/O is sent on the path unless the xref:{relative_path}storage-
path-resume.html[storage path resume] command is issued to continue I/O.

== Parameters

``-node {<nodename>|local}`` - Node name::
The name of the clustered node for which information is being displayed.

``-initiator <initiator name>`` - Initiator Port::
Initiator port that the clustered node uses.

``-target-wwpn <wwpn name>`` - Target Port::
Target World Wide Port Name. Port on the storage array that is being used.

{ ``[-lun-number <integer>]`` - LUN Number::
Logical Unit number. The range is: [0...65535]. If this parameter is not
specified, ONTAP resumes the entire path to an array.

| ``[-path-failure-threshold <integer>]`` - Max Number of Path Failures
Acceptable During wait-duration::
The path failure count, exceeding this value within wait duration will
quiesce the path.

``[-wait-duration <integer>]`` - Wait Duration in minutes }::
The time duration(minutes) in which path is monitored for path failures.

== Examples

The following example suspends I/O between node vbv3170f1b, port 0a and
the array port 50001fe1500a8669, LUN 1.
```

node::> storage path quiesce -node vbv3170f1b -initiator 0a -target-wwpn 50001fe1500a8669 -lun-number 1

```
The following example suspends I/O immediately between node vbv3170f1b,
port 0a and the array port 50001fe1500a8669.
```

node::> storage path quiesce -node vbv3170f1b -initiator 0a -target-wwpn 50001fe1500a8669

```
The following example suspends I/O between node vbv3170f1b, port 0a and
the array port 50001fe1500a8669 after reaching 10 or more errors in
duration of 5 mins.
```

node::> storage path quiesce -node vbv3170f1b -initiator 0a -target-wwpn 50001fe1500a8669 -path-failure -threshold 10 -wait-duration 5

```
== Related Links

* xref:{relative_path}storage-path-resume.html[storage path resume]




[[ID90a4cb8bd09fd41e9591d1746f317786]]
= storage path resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume I/O on a path to array

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage path resume` command continues I/O flow to an array LUN on a
path or the entire path that was previously quiesced. It also disables the
path failures monitoring feature, if it was enabled using the
xref:{relative_path}storage-path-quiesce.html[storage path quiesce]`-path-
failure-threshold```_count_`` command.

== Parameters

``-node {<nodename>|local}`` - Node name::
```

```
The name of the clustered node for which information is being displayed.

``-initiator <initiator name>`` - Initiator Port::
Initiator port that the clustered node uses.

``-target-wwpn <wwpn name>`` - Target Port::
Target World Wide Port Name. Port on the storage array that is being used.

``[-lun-number <integer>]`` - LUN Number::
Logical Unit number. The range is: [0...65535]. If this parameter is not
specified, ONTAP resumes the entire path to an array.

== Examples

The following example resumes I/O between node vbv3170f1b, port 0a and the
array port 50001fe1500a8669, LUN 1
```

node::> storage path resume -node vbv3170f1b -initiator 0a -target-wwpn 50001fe1500a8669 -lun-number 1

```
The following example resumes I/O between node vbv3170f1b, port 0a and the
array port 50001fe1500a8669
```

node::> storage path resume -node vbv3170f1b -initiator 0a -target-wwpn 50001fe1500a8669

```
== Related Links

* xref:{relative_path}storage-path-quiesce.html[storage path quiesce]




[[ID9394e5994e94745daee679028284f83d]]
= storage path show-by-initiator
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of paths to attached arrays from the initiator's
perspective

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `storage path show-by-initiator` command displays path based
statistics. The output is similar to the xref:{relative_path}storage-path-
show.html[storage path show] command but the output is listed by
initiator.

== 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>|local}]`` - Controller name::
The name of the clustered node for which information is being displayed.

``[-initiator <text>]`` - Initiator Port::
Initiator port that the clustered node uses.

``[-target-wwpn <text>]`` - Target Port::
Target World Wide Port Name. Port on the storage array that is being used.

``[-initiator-side-switch-port <text>]`` - Initiator Side Switch Port::
Switch port connected to the clustered node.

``[-target-side-switch-port <text>]`` - Target Side Switch Port::
Switch port connected to the array.

``[-array-name <array name>]`` - Array Name::
Name of the storage array that is connected to the cluster.

``[-tpgn <integer>]`` - Target Port Group Number::
TPGN refers to the target port group to which the target port belongs. A
target port group is a set of target ports which share the same LUN access
characteristics and failover behaviors.

``[-port-speed <text>]`` - Port Speed::
Port Speed of the specified port.

``[-path-io-kbps <integer>]`` - Kbytes of I/O per second on Path (Rolling
Average)::
Rolling average of I/O per second on the path.
```

```
``[-path-iops <integer>]`` - Number of IOPS on Path (Rolling Average)::
Rolling average of Kbytes of I/O per second on the path

``[-initiator-io-kbps <integer>]`` - Kbytes of I/O per second on Initiator
(Rolling Average)::
Rolling average of I/O per second on the initiator port.

``[-initiator-iops <integer>]`` - Number of IOPS on Initiator (Rolling
Average)::
Rolling average of Kbytes of I/O per second on the initiator port.

``[-target-io-kbps <integer>]`` - Kbytes of I/O per second to Target
(Rolling Average)::
Rolling average of I/O per second on the target port.

``[-target-iops <integer>]`` - Number of IOPS to Target (Rolling
Average)::
Rolling average of Kbytes of I/O per second on the target port.

== Examples
```

vnv3070f20b::> storage path show-by-initiator
Node: vnv3070f20b
Initiator I/O Initiator Side Path I/O Target Side Target I/O
Initiator (KB/s) Switch Port (KB/s) Switch Port (KB/s) Target Port Array Name
--------- ------------ ------------------- ----------- ------------------- ----------- --------------- ----------------
0a 3 vnbr3850s4:4 3 vnbr3850s5:15 3 200600a0b819e16f IBM_1722_1
0 vnbr3850s5:12 0 50060e80004291c0 HITACHI_DF600F_1
0c 35 vnci9124s54:1-6 35 vnci9124s54:1-24 35 200700a0b819e16f IBM_1722_1
0 vnci9124s54:1-22 0 50060e80004291c2 HITACHI_DF600F_1
4 entries were displayed.

```
== Related Links

* xref:{relative_path}storage-path-show.html[storage path show]




[[ID83deca55f05b8a4fec7856ad77d364fa]]
= storage path show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display a list of paths to attached arrays.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage path show` command displays path based statistics. The
default command shows:

* Node name
* Initiator port
* Target port
* Target IQN
* TPGN
* Port speeds
* Path I/O in Kbytes/sec
* IOPs



== 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.

| ``[-array  ]``::
Using this option displays:
+
  * Array name
  * Target port
  * Target IQN
  * Target I/O in Kbytes/sec
  * Target side switch port
  * Path I/O in Kbytes/sec
  * Initiator side switch port
  * Initiator I/O in Kbytes/sec
  * Initiator port



| ``[-by-target  ]``::
Using this option displays the same information as the array option, but
grouped by target port.
```

```
| ``[-detail  ]``::
Using this option displays the same information as the array and by-target
options, but adds the following:
+
  * Target IOPs
  * Target LUNs
  * Path IOPs
  * Path errors
  * Path quality
  * Path LUNs
  * Initiator IOPs
  * Initiator LUNs



| ``[-switch  ]``::
Using this option adds switch port information to the default display.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Controller name::
The name of the clustered node for which information is being displayed.

``[-array-name <array name>]`` - Array Name::
Name of the storage array that is connected to the cluster.

``[-target-wwpn <text>]`` - Target Port::
Target World Wide Port Name. Port on the storage array that is being used.

``[-initiator <text>]`` - Initiator Port::
Initiator port that the clustered node uses.

``[-initiator-side-switch-port <text>]`` - Initiator Side Switch Port::
Switch port connected to the clustered node.

``[-tpgn <integer>]`` - Target Port Group Number::
TPGN refers to the target port group to which the target port belongs. A
target port group is a set of target ports which share the same LUN access
characteristics and failover behaviors.

``[-port-speed <text>]`` - Port Speed::
Port Speed of the specified port.

``[-path-io-kbps <integer>]`` - Kbytes of I/O per second on Path (Rolling
```

```
Average)::
Rolling average of I/O per second on the path.

``[-path-iops <integer>]`` - Number of IOPS on Path (Rolling Average)::
Rolling average of Kbytes of I/O per second on the path

``[-initiator-io-kbps <integer>]`` - Kbytes of I/O per second on Initiator
(Rolling Average)::
Rolling average of I/O per second on the initiator port.

``[-initiator-iops <integer>]`` - Number of IOPS on Initiator (Rolling
Average)::
Rolling average of Kbytes of I/O per second on the initiator port.

``[-target-io-kbps <integer>]`` - Kbytes of I/O per second to Target
(Rolling Average)::
Rolling average of I/O per second on the target port.

``[-target-iops <integer>]`` - Number of IOPS to Target (Rolling
Average)::
Rolling average of Kbytes of I/O per second on the target port.

``[-target-side-switch-port <text>]`` - Target Side Switch Port::
Switch port connected to the array.

``[-path-link-errors <integer>]`` - Link Error count on path::
Fibre Channel link error count.

``[-path-quality <integer>]`` - Percentage of weighted error threshold::
A number representing the threshold of errors that is allowed on the path.
Path quality is a weighted error value. When the error weight of a path
exceeds the threshold, I/O is routed to a different path.

``[-path-lun-in-use-count <integer>]`` - Number of LUNs in the in-use
state on this path::
Number of LUNs on this path.

``[-initiator-lun-in-use-count <integer>]`` - Number of LUNs in the in-use
state on this initiator::
Number of LUNs on this initiator.

``[-target-lun-in-use-count <integer>]`` - Number of LUNs in the in-use
state on this target::
Number of LUNs on this target.

``[-vmdisk-device-id <integer>]`` - Virtual disk device ID::
Common device identifier, shared by a VM and its hypervisor, of a virtual
```

```
disk. On ESX servers, this is the Disk ID component of a virtual device
node, with a value of 0 to 15.

``[-path-failure-threshold <integer>]`` - Max number of path failures
acceptable in wait-duration::
The path failure count, exceeding this value within wait duration will
quiesce the path.

``[-wait-duration <integer>]`` - Wait Duration in minutes::
The time duration(minutes) in which path is monitored for path failures.

== Examples

The following example shows the default display.
```

vbv3170f2a::> storage path show
Path I/O
Node Initiator Array Target Port TPGN Speed (KB/s) IOPS
--- --- --- --- --- --- ---
vbv3170f2a-01 0b 50001fe1500a866c 2 2 Gb/S 6 2
vbv3170f2a-01 0b 50001fe1500a866d 2 2 Gb/S 0 0
vbv3170f2a-01 0c 50001fe1500a866e 4 4 Gb/S 0 0
vbv3170f2b-03 0a 50001fe1500a866d 1 2 Gb/S 3 1
vbv3170f2b-03 0c 50001fe1500a866f 4 4 Gb/S 3 1
5 entries were displayed.

```
The following example shows how the information is displayed with the
array option.
```

vnv3070f20b::> storage path show -array
Node: vnv3070f20b
Target I/O Target Side Path I/O Initiator Side Initiator I/O Initiator
Array Name Target Port (KB/s) Switch Port (KB/s) Switch Port (KB/s) Port
--- --- --- --- --- --- --- ---
HITACHI_DF600F_1 50060e80004291c0 0 vnbr3850s5:12 0 vnbr3850s4:4 3 0a
50060e80004291c2 0 vnci9124s54:1-22 0 vnci9124s54:1-6 26 0c
IBM_1722_1 200600a0b819e16f 3 vnbr3850s5:15 3 vnbr3850s4:4 3 0a
200700a0b819e16f 26 vnci9124s54:1-24 26 vnci9124s54:1-6 26 0c
4 entries were displayed.

```
The following example shows how the information is displayed when grouped
by target.
```

vnv3070f20b::> storage path show -by-target
Node: vnv3070f20b
Array Name: HITACHI_DF600F_1
Target I/O Target Side Path I/0 Initiator Side Initiator I/O Initiator

Target Port (KB/s) Switch Port (KB/s) Switch Port (KB/s) Port
---------- ---------- ---------- ---------- ---------- ---------- ---------
50060e80004291c0 0 vnbr3850s5:12 0 vnbr3850s4:4 3 0a
50060e80004291c2 0 vnci9124s54:1-22 0 vnci9124s54:1-6 26 0c
Node: vnv3070f20b
Array Name: IBM_1722_1
Target I/O Target Side Path I/0 Initiator Side Initiator I/O Initiator
Target Port (KB/s) Switch Port (KB/s) Switch Port (KB/s) Port
---------- ---------- ---------- ---------- ---------- ---------- ---------
200600a0b819e16f 3 vnbr3850s5:15 3 vnbr3850s4:4 3 0a
200700a0b819e16f 26 vnci9124s54:1-24 26 vnci9124s54:1-6 26 0c
4 entries were displayed.

```
The following example shows how the information is displayed with the
switch option.
```

vbv3170f2b::> storage path show -switch
Target Side Initiator Side Path I/O
Node Initiator Array Target Port Switch Port Switch Port TPGN Speed (KB/s) IOPS
---------- --------- ---------- ---------- ---------- ------ ------- ---------- ----------
vbv3170f2a-01 0b 50001fe1500a866c vbbr300s1:6 vbbr300s1:2 2 2 Gb/S 9 3
vbv3170f2a-01 0b 50001fe1500a866d vbbr300s1:7 vbbr300s1:2 2 2 Gb/S 0 0
vbv3170f2a-01 0c 50001fe1500a866e vbci9124s2:1-7 vbci9124s2:1-3 4 4 Gb/S 0 0
vbv3170f2b-03 0a 50001fe1500a866d vbbr300s1:7 vbbr300s1:3 1 2 Gb/S 4 1
vbv3170f2b-03 0c 50001fe1500a866f vbci9124s2:1-8 vbci9124s2:1-4 4 4 Gb/S 4 1
5 entries were displayed.

```
:leveloffset: -1


= storage pool commands

:leveloffset: +1


[[IDbec6b139e24bd09807582b721cab6422]]
= storage pool add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add disks to a storage pool

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `storage pool add` command increases the total capacity of an existing
storage pool by adding the specified SSDs to the storage pool. The disks
are split into four equal partitions and added to each of the allocation
units of the storage pool. If any allocation units from the storage pool
have already been allocated to an aggregate, the cache capacity of that
aggregate is increased automatically.

For example, if an SSD with a usable size of 745 GB is added to a storage
pool that is part of four aggregates, each aggregate will grow its cache
capacity by 186.2 GB. If a different allocation is desired, create a new
storage pool using the xref:{relative_path}storage-pool-
create.html[storage pool create] command.

NOTE: Storage pools are not supported in all-flash aggregates.



== Parameters

``-storage-pool <storage pool name>`` - Storage Pool Name::
This parameter specifies the storage pool to which disks are to be added.

{ ``[-disk-count <integer>]`` - Number of Disks to Add in Storage Pool::
This parameter specifies the number of disks that are to be added to the
storage pool. The disks to be added come from the pool of spare disks.

``[-nodes {<nodename>|local}]`` - Nodes From Which Spares Should be
Selected::
This parameter specifies a list of nodes from which SSD disks are selected
for addition to the storage pool. If this parameter is not specified,
disks to be added to the storage pool can be selected from both the nodes
sharing the storage pool. Use this parameter to restrict the selection of
spare disks to one particular node.

| ``[-disk-list <disk path name>,...]`` - List of Spare Disks }::
This parameter specifies a list of disks to be added to the storage pool.
In an HA configuration, SSDs being added to a storage pool can be owned by
either node in the HA pair.

{ ``[-quiet <true>]`` - Confirmations off [.advaced-priv]#(privilege:
advanced)#::
When set to ``_true_`` , this parameter specifies the operation should be
executed without pausing for confirmation.

``[-simulate <true>]`` - Simulate Storage Pool Addition::
```

```
When set to ``_true_`` , this parameter specifies the operation should be
performed as a simulation. The command reports which aggregates would grow
automatically as a result of adding the disks to the storage pool. The
disks are not added to the storage pool.

== Examples

In this example, the user requests a report detailing the changes that
would occur if a new disk is added to the storage pool ``_SP1_`` . In this
case, 186.2 GB of cache is added to the Flash Pool aggregates
nodeA_flashpool_1 and nodeB_flashpool_1. There are two unprovisioned
allocation units in the storage pool and therefore the storage pool
available capacity also grows by 372.5 GB.
```

cluster1::> storage pool add -storage-pool SP1 -disk-list 1.0.23 -simulate

This operation will result in capacity being allocated in the following way:
Container Capacity Current New
Name To Be Added Size Size
---------- --- --- ---
nodeA_flashpool_1 186.2GB 558.7GB 744.9GB
nodeB_flashpool_1 186.2GB 558.7GB 744.9GB
(Available Capacity) 372.5GB 1.09TB 1.45TB

```
The following example adds one disk to a storage pool named ``_SP1_`` .
The spare disks are selected from either local node or its partner or both
based on spare availability.
```

cluster-1::> storage pool add -storage-pool SP1 -disk-count 1

Info: The following disks will be added to storage pool "SP1":
Disk Size Type Owner
---------- --- --- ---
1.0.12 744.9GB SSD cluster-1-01
New Allocation Unit Size: 744.8GB
Capacity will be allocated in the following way:
Container Capacity Current New
Name To Be Added Size Size
---------- --- --- ---
nodeA_flashpool_1 186.2GB 558.7GB 744.9GB
nodeB_flashpool_1 186.2GB 558.7GB 744.9GB
(Available Capacity) 372.5GB 1.09TB 1.45TB
Are you sure you want to continue with this operation?
{y|n}: y
[Job 48] Job succeeded: storage pool add job for "SP1" completed successfully

```
== Related Links
```

```
* xref:{relative_path}storage-pool-create.html[storage pool create]




[[IDab8c872b4f6bbd5a6045a2b748abd6e4]]
= storage pool create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new storage pool

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage pool create` command creates an SSD storage pool using a
given list of spare SSDs.

When a storage pool is created, ONTAP splits the capacity provided by the
SSDs into four equally-sized allocation units. In an HA configuration, two
allocation units (containing 50% of the total capacity) are assigned to
each node in the HA pair. This assignment can be modified using the
xref:{relative_path}storage-pool-reassign.html[storage pool reassign]
command.

After the storage pool is created, its allocation units can be provisioned
into Flash Pool aggregates using the xref:{relative_path}storage-
aggregate-add-disks.html[storage aggregate add-disks] command and the `-
storage-pool` parameter.

NOTE: Storage pools are not supported in All-Flash aggregates.



== Parameters

``-storage-pool <storage pool name>`` - Storage Pool Name::
This parameter specifies the name of the storage pool that is to be
created. The SSDs are partitioned and placed into the new storage pool.

{ ``[-nodes {<nodename>|local}]`` - Nodes Sharing the Storage Pool::
This parameter specifies a list of nodes from which SSD disks are selected
```

```
to create the storage pool. If two nodes are specified then they need to
be in HA configuration. Spare disks are selected from either node or its
partner or both. If this parameter is not specified, storage pool will be
created by selecting disks from either the node or its partner or both
from where command is run.

``-disk-count <integer>`` - Number of Disks in Storage Pool::
This parameter specifies the number of disks that are to be included in
the storage pool. The disks in this newly created storage pool come from
the pool of spare disks. The smallest disks in this pool are added to the
storage pool first, unless you specify the `-disk-size` parameter.

``[-disk-size {<integer>[KB|MB|GB|TB|PB]}]`` - Disk Size::
This parameter specifies the size of the disks on which the storage pool
is to be created. Disks with a usable size between 95% and 105% of the
specified size are selected.

| ``-disk-list <disk path name>,...`` - Disk List for Storage Pool
Creation }::
This parameter specifies a list of SSDs to be included in the new storage
pool. The SSDs must be spare disks and can be owned by either node in an
HA pair.

``[-simulate <true>]`` - Simulate Storage Pool Creation::
This option simulates the storage pool creation and prints the allocation
unit size that would be used for the storage pool.

== Examples

The following example creates a storage pool named SP1. The storage pool
contains 3 SSD disks, the spare disks selected are from either local node,
or its partner or both based on spare availability.
```

cluster1::> storage pool create -storage-pool SP1 -disk-count 3

```
The following example creates a storage pool named SP2. The storage pool
contains 3 SSD disks, the spare disks selected are from either node0, or
its partner node1 or both based on spare availability.
```

cluster1::> storage pool create -storage-pool SP2 -disk-count 3 -nodes node0,node1

```
The following example creates a storage pool named SP3 from four SSDs
using disk list.
```

cluster1::> storage pool create -storage-pool SP3 -disk-list 1.0.13, 1.0.15, 1.0.17, 1.0.19

```
== Related Links

* xref:{relative_path}storage-pool-reassign.html[storage pool reassign]
* xref:{relative_path}storage-aggregate-add-disks.html[storage aggregate
add-disks]




[[IDac7b0660fc669da628c65dfcb3fd12c8]]
= storage pool delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing storage pool

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage pool delete` command deletes an existing SSD storage pool. At
the end of the operation, the SSDs are converted back to spare disks.

== Parameters

``-storage-pool <storage pool name>`` - Storage Pool Name::
This parameter specifies the storage pool that you want to delete. You can
delete the storage pool only if all of the allocation units in the storage
pool are available.

== Examples

Verify that storage pool ``_SP3_`` is ready for deletion by confirming it
has four available allocation units and then delete it.
```

cluster1::> storage pool show-available-capacity -storage-pool SP3
Storage SyncMirror Allocation Unit Total
Node Storage Pool Type Pool Unit size Count Usable Size
--- --- --- --- --- --- ---
node-a SP3 SSD Pool0 372.5GB 2 744.9GB
node-b SP3 SSD Pool0 372.5GB 2 744.9GB

```
2 entries were displayed.

cluster1::> storage pool delete -storage-pool SP3

Warning: Are you sure you want to delete storage pool "SP3"? {y|n}: y
[Job 313] Job succeeded: storage pool delete job for "SP3" completed successfully
```

```
[[ID4286b0b445ec7fdc885cded38c1a46cb]]
= storage pool reassign
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reassign capacity from one node to another node in storage pool

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage pool reassign` command changes the ownership of unprovisioned
(available) storage pool allocation units from one HA partner to the other
for an existing storage pool.

== Parameters

``-storage-pool <storage pool name>`` - Storage Pool Name::
This parameter specifies the storage pool within which available capacity
is reassigned from one node to another.

``-from-node {<nodename>|local}`` - Reassign Available Capacity from This
Node::
This parameter specifies the name of the node that currently owns the
allocation units.

``-to-node {<nodename>|local}`` - Reassign Available Capacity to This
Node::
This parameter specifies the name of the node that will now own the
allocation units.

``-allocation-units <integer>`` - Allocation Units::
This parameter specifies the number of allocation units to be reassigned.

== Examples

Move an available allocation unit from node-b to node-a in preparation for
provisioning the allocation units on node-a.
```

cluster1::*> storage pool show-available-capacity -storage-pool SP2
Storage SyncMirror Allocation Unit Total
Node Storage Pool Type Pool Unit size Count Usable Size

---------- --------------- ------- ---------- ---------- ----- -----------
node-a SP2 SSD Pool0 744.9GB 2 1.45TB
node-b SP2 SSD Pool0 744.9GB 1 744.9GB
2 entries were displayed.

cluster1::*> storage pool reassign -storage-pool SP2 -from-node node-b -to-node node-a -allocation-units 1
[Job 310] Job succeeded: storage pool reassign job for "SP2" completed successfully

cluster1::*> storage pool show-available-capacity -storage-pool SP2
Storage SyncMirror Allocation Unit Total
Node Storage Pool Type Pool Unit size Count Usable Size
---------- --------------- ------- ---------- ---------- ----- -----------
node-a SP2 SSD Pool0 744.9GB 3 2.18TB
node-b SP2 SSD Pool0 744.9GB 0 0B
2 entries were displayed.

```
[[IDd29b302380985daf47d5721a6e59aadf]]
= storage pool rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename storage pool

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage pool rename` command changes the name of the storage pool.

== Parameters

``-storage-pool <storage pool name>`` - Storage Pool Name::
This parameter specifies the storage pool name.

``-new-name <storage pool name>`` - New Name of the Storage Pool::
This parameter specifies the new name of the storage pool.

== Examples

Renaming storage pool "sp1" to "sp2"
```

cluster-1::> storage pool show
Storage
Storage Pool Type #Disks Nodes Total Size

---
sp1 SSD 4 node-a, 10.44GB
node-b

cluster-1::> storage pool rename -storage-pool sp1 -new-name sp2

cluster-1::> storage pool show
Storage
Storage Pool Type #Disks Nodes Total Size
---
sp2 SSD 4 node-a, 10.44GB
node-b

```
[[ID657a4e0e8bc1f4ca6637e274a9543110]]
= storage pool show-aggregate
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display aggregates provisioned from storage pools

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage pool show-aggregate` command displays allocation information
for SSD storage pools in the cluster. The command output depends upon the
parameter or parameters specified with the command. If no parameters are
specified, the command displays information about allocations of all
storage pools in the cluster.

== 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.

``[-storage-pool <storage pool name>]`` - Name of Storage Pool::
Selects the storage pools that match this parameter value.
```

```
``[-aggregate <aggregate name>]`` - Aggregate::
Selects the storage pools that match this parameter value.

``[-capacity {<integer>[KB|MB|GB|TB|PB]}]`` - Capacity::
Selects the storage pools that match this parameter value.
+
Capacity includes space provided by data and parity portions of each
allocation unit. Only the data portions of each allocation unit contribute
to the cache or usable capacity of Flash Pool.

``[-allocated-unit-count <integer>]`` - Number of AU's Assigned to This
Aggregate::
Selects the storage pools that match this parameter value.

``[-original-owner <text>]`` - Original Owner Name::
Selects the storage pools that match this parameter value.

``[-node {<nodename>|local}]`` - Node::
Selects the storage pools that match this parameter value.

== Examples

Display information about the aggregate or aggregates using a storage pool
called ``_SP2_`` :
```

cluster1::> storage pool show-aggregate -storage-pool SP2 -instance
Name of Storage Pool: SP2
Aggregate: node2_flashpool_1
Capacity: 744.9GB
Number of AU’s Assigned to This Aggregate: 1
Original Owner Name: node2
Node: node2

```
[[ID12b50efac5e9a19b3e73cf6328a9442c]]
= storage pool show-available-capacity
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display available capacity of storage pools

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `storage pool show-available-capacity` command displays information
about available capacity in SSD storage pools on each node in the cluster.
The command output depends upon the parameter or parameters specified with
the command. If no parameters are specified, the command displays
information about available capacities in all shared pools in the cluster.

Storage pool available capacity is data storage space that has not yet
been provisioned into Flash Pool. Allocation units might be provisioned
into aggregates using the xref:{relative_path}storage-aggregate-add-
disks.html[storage aggregate add-disks] command and the `-storage-pool`
parameter.

NOTE: All storage pool available capacity can be provisioned into
aggregates. Available capacity within a storage pool is not used to
protect against a disk failure. In the case of an SSD failure or predicted
failure, ONTAP moves a suitable whole SSD spare disk from outside the
storage pool into the storage pool and begins the recovery process (using
either reconstruction or Rapid RAID Recovery, whichever is appropriate).



== 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.

``[-storage-pool <storage pool name>]`` - Name of Storage Pool::
Selects the available capacities that match this parameter value.

``[-node {<nodename>|local}]`` - Node::
Selects the available capacities that match this parameter value.

``[-allocation-unit-size {<integer>[KB|MB|GB|TB|PB]}]`` - Allocation Unit
Size::
Selects the available capacities that match this parameter value.
+
Allocation units are the units of storage capacity that are available to
be provisioned into aggregates.
```

```
``[-storage-type <SSD>]`` - Type of Storage Pool::
Selects the available capacities that match this parameter value. Only the
``_SSD_`` type is supported for this version of ONTAP.

``[-allocation-unit-count <integer>]`` - Number of Allocation Units
Available::
Selects the available capacities that match this parameter value.
+
Allocation units are the units of storage capacity that are available to
be provisioned into aggregates. Each allocation unit is one minimum unit
of allocation (MUA) and its capacity is given as `allocation-unit-size` .

``[-syncmirror-pool <text>]`` - SyncMirror Pool::
Selects the available capacities that match this parameter value.
+
The SyncMirror pool of an allocation unit must match the SyncMirror pool
of the disks of the aggregate when adding allocation units into an
aggregate.
+
Mirroring of aggregates that are provisioned from SSD storage pools is not
supported.

``[-available-size {<integer>[KB|MB|GB|TB|PB]}]`` - Total Usable Available
Size::
Selects the available capacities that match this parameter value.
+
The `available-size` is the sum of the capacities of the allocation units
that are assigned but not yet provisioned. The amount of `available-size`
that is contributed to the cache or usable capacity of an aggregate
depends upon the RAID type used when provisioning the allocation units.

== Examples

In this example, two nodes of an HA pair share available capacity from two
storage pools, ``_SP1_`` and ``_SP2_`` . There are a total of 5 allocation
units that have not yet been provisioned.
```

cluster1::> storage pool show-available-capacity
Storage SyncMirror Allocation Unit Total
Node Storage Pool Type Pool Unit size Count Usable Size
--- --- --- --- --- --- ---
node-a SP1 SSD Pool0 558.7GB 1 558.7GB
node-b SP1 SSD Pool0 558.7GB 1 558.7GB
node-a SP2 SSD Pool0 744.9GB 2 1.45TB
node-b SP2 SSD Pool0 744.9GB 1 744.9GB

```
== Related Links

* xref:{relative_path}storage-aggregate-add-disks.html[storage aggregate
add-disks]




[[IDfcbc2508fe0b7a27e368822f3060ca6e]]
= storage pool show-disks
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display disks in storage pools

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage pool show-disks` command displays information about disks in
storage pools in the cluster. The command output depends on the parameter
or parameters specified with the command. If no parameters are specified,
the command displays information about all disks in all storage pools in
the cluster.

== 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.

``[-storage-pool <storage pool name>]`` - Name of Storage Pool::
Selects the storage pools that match this parameter value.

``[-disk <disk path name>]`` - Name of the disk::
Selects the storage pools with the disks that match this parameter value.

``[-disk-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS  | SSD |
```

```
VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}]`` - Disk Type::
Selects the storage pools with the disks that match this parameter value.
Only the ``_SSD_`` type is supported for this version of ONTAP.

``[-usable-size {<integer>[KB|MB|GB|TB|PB]}]`` - Disk Usable Size::
Selects the storage pools with the disks that match this parameter value.
+
In this command, `usable-size` refers to the sum of the capacities of all
of the partitions on the disk.

``[-total-size {<integer>[KB|MB|GB|TB|PB]}]`` - Total Size::
Selects the storage pools with the disks that match this parameter value.

``[-node-list <nodename>,...]`` - List of Nodes::
Selects the storage pools with the disks that are visible to all of the
specified nodes.

== Examples

Show information about SSDs in a storage pool called ``_SP2_`` .
```

cluster1::> storage pool show-disks -storage-pool SP2

Storage Pool Name: SP2
Storage
Disk Type Usable Size Total Size
-------- ------- ---------- ----------
1.0.16 SSD 745.0GB 745.2GB
1.0.18 SSD 745.0GB 745.2GB
1.0.20 SSD 745.0GB 745.2GB
1.0.22 SSD 745.0GB 745.2GB

```
[[IDd9bbc304b1d3120843118b9d6da8acc9]]
= storage pool show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display details of storage pools

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `storage pool show` command displays information about SSD storage
pools in the cluster. By default, the command displays information about
all storage pools in the cluster. You can specify parameters to limit the
output to a specific set of storage pools.

== 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.

``[-storage-pool <storage pool name>]`` - Storage Pool Name::
Selects the storage pools that match this parameter value.

``[-storage-pool-uuid <UUID>]`` - UUID of Storage Pool::
Selects the storage pools that match this parameter value.

``[-nodes {<nodename>|local}]`` - Nodes Sharing the Storage Pool::
Selects the storage pools that match this parameter value.
+
In an HA pair, either node name may be specified.

``[-disk-count <integer>]`` - Number of Disks in Storage Pool::
Selects the storage pools that match this parameter value.

``[-allocation-unit-size {<integer>[KB|MB|GB|TB|PB]}]`` - Allocation Unit
Size::
Selects the storage pools that match this parameter value.
+
Allocation units represent the unit of storage allocated to aggregates
from this storage pool.

``[-allocation-unit-data-size-raid4 {<integer>[KB|MB|GB|TB|PB]}]`` -
Allocation Unit Data Size for RAID4 and RAID-EP::
This parameter shows the amount of additional data capacity provided if an
allocation unit from this storage pool was added to an aggregate with `-
raidtype` as ``_raid4_`` .

``[-allocation-unit-data-size-raid-dp {<integer>[KB|MB|GB|TB|PB]}]`` -
Allocation Unit Data Size for RAID-DP::
This parameter shows the amount of additional data capacity provided if an
allocation unit from this storage pool was added to an aggregate with `-
```

```
raidtype` as ``_raid_dp_`` .

``[-allocation-unit-data-size-raid-tec {<integer>[KB|MB|GB|TB|PB]}]`` -
Allocation Unit Data Size for RAID-TEC::
This parameter shows the amount of additional data capacity provided if an
allocation unit from this storage pool was added to an aggregate with `-
raidtype` as ``_raid_tec_`` .

``[-storage-type <SSD>]`` - Storage Type::
Selects the storage pools that match this parameter value.
+
Only the `SSD` type is supported for this version of ONTAP.

``[-pool-usable-size {<integer>[KB|MB|GB|TB|PB]}]`` - Storage Pool Usable
Size::
Selects the storage pools that match this parameter value.
+
The `pool-usable-size` is the sum of the capacities of the allocation
units that are assigned to nodes but not yet provisioned. The amount of
`pool-usable-size` that is contributed to the cache or usable capacity of
an aggregate depends upon the RAID type used when provisioning the
allocation units.

``[-pool-total-size {<integer>[KB|MB|GB|TB|PB]}]`` - Storage Pool Total
Size::
Selects the storage pools that match this parameter value.
+
The `pool-total-size` is the sum of the capacities of allocation units
belonging to this storage pool.

``[-is-healthy {true|false}]`` - Is Pool Healthy?::
Selects the storage pools that match this parameter value.
+
For storage pools with `is-healthy```_false_`` , the `unhealthy-reason`
parameter provides more information.
+
`is-healthy` must be ``_true_`` to provision allocation units from a
storage pool into an aggregate.

``[-pool-state <State of the Storage Pool>]`` - State of the Storage
Pool::
Selects the storage pools that match this parameter value. Possible states
are:
+
  * normal - the storage pool is operating normally.
  * degraded - the storage pool has one or more failed disks.
```

```
  * creating - the storage pool is being created.
  * deleting - the storage pool is being deleted.
  * reassigning - allocation units are being reassigned from one node to
another.
  * growing - allocation units in the storage pool are expanding due to
the addition of new capacity into the storage pool.



``[-unhealthy-reason <text>]`` - Reason for Storage Pool Being Unhealthy::
Selects the storage pools that match this parameter value.
+
The message provided gives additional details about why the storage pool
is unhealthy.

``[-current-operation-job-id <integer>]`` - Job ID of the Currently
Running Operation::
Selects the storage pools that match this parameter value.
+
Long-running operations associated with storage pools will be managed via
jobs. For example, if you provision allocation units from a storage pool
into an aggregate and the disks associated with the storage pool need to
be zeroed, the operation will be completed via a job.

== Examples

Display the storage pools in the cluster.
```

cluster1::> storage pool show
Storage Pool Type #Disks Nodes Total Size
--- --- --- --- ---
LargeSP SSD 10 noda-a,node-b 7.27TB
SmallSP SSD 2 noda-a,node-b 1.45TB
2 entries were displayed.

```
The following example displays the details of a storage pool named
SmallSP. Only one of its four allocation unit has been provisioned, so 75%
of its size is available (usable).
```

cluster1::> storage pool show -storage-pool SmallSP
Storage Pool Name: SmallSP
UUID of Storage Pool: 60f2f1b9-e60f-11e3-a5e7-00a0981899a2
Nodes Sharing the Storage Pool: node-a, node-b
Number of Disks in Storage Pool: 2
Allocation Unit Size: 372.5GB
Storage Type: SSD
Storage Pool Usable Size: 1.09TB

Storage Pool Total Size: 1.45TB
Is Pool Healthy?: true
State of the Storage Pool: normal
Reason for storage pool being unhealthy: -
Job ID of the Currently Running Operation: -

```
:leveloffset: -1


= storage port commands

:leveloffset: +1


[[ID25fbd7b92e327600a3a6a3d5c176461b]]
= storage port disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable a storage port

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage port disable` command disables a specified storage port.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node on which the port resides.

``-port <text>`` - Port::
Use this parameter to specify the port that needs to be disabled.

``[-force <true>]`` - Force [.advaced-priv]#(privilege: advanced)#::
Use this optional parameter to force the disabling of the storage port.
The parameter can be used to disable the specified port even if some
devices can only be accessed using this port. Note that doing so might
cause multiple device failures.

== Examples

The following example disables port 0a on node node1:
```

cluster1::> storage port disable -node node1 -port 0a

```
[[IDeb119ff86b40c3101edbc2ccb7315767]]
= storage port enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable a storage port

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage port enable` command enables a specified storage port.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node on which the port resides.

``-port <text>`` - Port::
Use this parameter to specify the port that needs to be enabled.

== Examples

The following example enables port 0a on node node1:
```

cluster1::> storage port enable -node node1 -port 0a

```
[[IDddaadc790436ffe6a9eb02b321bde4af]]
= storage port modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a storage port

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage port modify` command modifies the attributes of a storage
port.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the name of the node.

``-port <text>`` - Port::
This parameter specifies the name of the port.

``[-mode {storage|network}]`` - Mode::
This parameter specifies the usage mode for an Ethernet port - storage or
network use. The mode parameter only applies to Ethernet ports that are
not dedicated to storage connectivity.

``[-force <true>]`` - Force [.advaced-priv]#(privilege: advanced)#::
This parameter allows the operation to override warnings.

== Examples

The following example places a nondedicated Ethernet port in storage mode:
```

cluster1::> storage port show
Speed VLAN
Node Port Type Mode (Gb/s) State Status ID
--- --- --- --- --- --- --- ---
cluster1-01
e0b ENET network - - - -
e0c ENET - 100 enabled online 30
e0e ENET storage 0 enabled offline 30

```
e0f ENET - 100 enabled online 30
cluster1-02
e0b ENET storage 0 enabled offline 30
e0c ENET - 100 enabled online 30
e0e ENET storage 0 enabled offline 30
e0f ENET - 100 enabled online 30
8 entries were displayed.

cluster1::> storage port modify -node cluster1-01 -port e0b -mode storage

cluster1::> storage port show
Speed VLAN
Node Port Type Mode (Gb/s) State Status ID
-------------- ---- ----- ------- ------ -------- --------------- ----
cluster1-01
e0b ENET storage 0 enabled offline 30
e0c ENET - 100 enabled online 30
e0e ENET storage 0 enabled offline 30
e0f ENET - 100 enabled online 30
cluster1-02
e0b ENET storage 0 enabled offline 30
e0c ENET - 100 enabled online 30
e0e ENET storage 0 enabled offline 30
e0f ENET - 100 enabled online 30
8 entries were displayed.
```

```
[[IDc4747ef8f69933b4e1e60cfec33e8682]]
= storage port rescan
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rescan a storage port

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage port rescan` command rescans a specified storage port. This
command is not supported on Ethernet storage ports (type = ENET).

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node on which the port resides.

``-port <text>`` - Port [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the port that needs to be rescanned.

== Examples

The following example rescans port 0a on node node1:
```

cluster1::> storage port rescan -node node1 -port 0a

```
[[ID2a3e31d2634d18fcb133f9badb451549]]
= storage port reset-device
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset a device behind a storage port

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

 The `storage port reset-device` command resets a device behind a port. If
the device is behind a SAS port, you need to specify the shelf name and
bay ID where the device resides. If the device is behind a FC port, you
need to specify the loop ID of the device. This command is not supported
on Ethernet storage ports (type = ENET).

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node on which the port resides.

``-port <text>`` - Port [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the port used to reset the device.

{ ``-shelf-name <text>`` - Shelf Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the shelf where the device resides.

``-bay-id <integer>`` - Bay ID [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the bay where the device resides.

| ``-loop-id <integer>`` - Loop ID [.advaced-priv]#(privilege: advanced)#
}::
Use this parameter to specify the loop ID of the device.

== Examples

The following example resets a device behind SAS port 0a on node node1:
```

cluster1::> storage port reset-device -node node1 -port 0a -shelf-name 1.0 -bay-id 10

```
The following example resets a device behind FC port 1b on node node1:
```

cluster1::> storage port reset-device -node node1 -port 1b -loop-id 20

```
[[ID9d52ae72cc5c89f141b61bd82a1da12a]]
= storage port reset
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset a storage port

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage port reset` command resets a specified storage port. This
command is not supported on Ethernet storage ports (type = ENET).

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node on which the port resides.

``-port <text>`` - Port [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the port that needs to be reset.

== Examples

The following example resets port 0a on node node1:
```

cluster1::> storage port reset -node node1 -port 0a

```
[[IDf88942ea1653671ec87acc2c8d16e6ff]]
= storage port show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show storage port information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage port show` command displays information about the storage
ports in the cluster. If no parameters are specified, the default command
displays the following information about the storage ports:

* Node
* Port
* Type
* Speed
* State
* Status

To display detailed profile information about a single storage port, use
the `-node` and `-port` parameters.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all the storage ports, in column style
output.

| ``[-errors  ]``::
Displays the following error status information about the storage ports
which have errors:
+
  * Error type
  * Error severity
  * Error description



| ``[-instance  ]`` }::
Displays expanded information about all the storage ports in the system.
If a storage port is specified, then this parameter displays detailed
information for that port only.

``[-node {<nodename>|local}]`` - Node::
Displays detailed information about the storage ports on the specified
```

```
node.

``[-port <text>]`` - Port::
Selects the ports with the specified port name.

``[-port-type {Unknown|SAS|FC|ENET}]`` - Port Type::
Selects the ports of the specified type.

``[-port-speed {0|1|1.5|2|2.5|3|4|5|6|8|10|12|14|16|25|32|40|64|100}]`` -
Port Speed::
Selects the ports with the specified speed.

``[-state {enabled|disabled|enable-pending|disable-pending}]`` - Port
State::
Selects the ports with the specified state.

``[-status {unknown|online|online-degraded|offline|link-down|transient-
mode}]`` - Port Status::
Selects the ports with the specified operational status.

``[-description <text>]`` - Description::
Selects the ports with the specified description.

``[-firmware-rev <text>]`` - Firmware Revision::
Selects the ports with the specified firmware revision.

``[-serial-number <text>]`` - Serial Number::
Selects the ports with the specified serial number.

``[-is-dedicated {true|false}]`` - Is Dedicated Storage Port?::
Selects the ports that match the specified value for storage-only ports.
This value is always true for FC and SAS ports, as well as for ENET ports
that are dedicated to storage.

``[-part-number <text>]`` - Part Number::
Selects the ports with the specified part number.

``[-connection-mode {Unknown|Loop|Point-to-point}]`` - Connection Mode::
Selects the ports with the specified connection mode.

``[-wwnn <FC WWN>]`` - World Wide Node Name::
Selects the ports with the specified World Wide Node Name.

``[-wwpn <FC WWN>]`` - World Wide Port Name::
Selects the ports with the specified World Wide Port Name.

``[-board-name <text>]`` - Board Name::
```

```
Selects the ports with the specified board name.

``[-connector-capabilities <integer>,...]`` - Connector Capabilities::
Selects the ports with the specified list of connector capabilities.

``[-wwn <FC WWN>]`` - Base World Wide Name::
Selects the ports with the specified World Wide Name.

``[-mfg-part-number <text>]`` - MFG Part Number::
Selects the ports with the specified manufacturer part number.

``[-nvdata-rev <text>]`` - NVDATA Revision::
Selects the ports with the specified NVDATA revision.

``[-date-code <text>]`` - Date Code::
Selects the ports with the specified date code.

``[-connector-technology {active-copper|passive-copper|optical}]`` -
Connector Technology::
Selects the ports with the specified connector technology.

``[-phy-id <integer>,...]`` - Phy ID::
Selects the ports that have phys with the specified phy ID.

``[-phy-state {enabled|disabled}]`` - Phy State::
Selects the ports that have phys with the specified state.

``[-phy-status {unknown|online|offline|speed-negotiation-failed|sata-oob-
failed}]`` - Phy Status::
Selects the ports that have phys with the specified status.

``[-phy-speed {0|1|1.5|2|2.5|3|4|5|6|8|10|12|14|16|25|32|40|64|100}]`` -
Phy Speed::
Selects the ports that have phys with the specified speed.

``[-mac-address <text>]`` - MAC Address::
Selects ports that match the specified MAC address.

``[-mode {storage|network}]`` - Mode::
Selects the ports with the specified mode.

``[-vlan-id <integer>]`` - VLAN ID::
Selects the ports with the specified VLAN ID.

``[-vendor-id <text>]`` - (DEPRECATED) Vendor ID::
+
NOTE: This parameter has been deprecated and may be removed in a future
```

```
version of ONTAP.
+
 Selects the ports with the specified vendor ID.

``[-vendor-part-id <text>]`` - (DEPRECATED) Vendor part ID::
+
NOTE: This parameter has been deprecated and may be removed in a future
version of ONTAP.
+
 Selects the ports with the specified vendor part ID.

``[-device-type <text>]`` - Device type::
Selects ports that match the specified device type.

``[-hw-rev <text>]`` - Hardware Revision::
Selects the ports with the specified hardware revision.

``[-error-type {unknown|online|online-degraded|offline|link-
down|transient-mode}]`` - Error Type::
Selects the ports with the specified error type.

``[-error-severity {unknown|notice|warning|error|critical}]`` - Error
Severity::
Selects the ports with the specified error severity.

``[-error-text <text>]`` - Error Text::
Selects the ports with the specified error text.

``[-corrective-action <text>]`` - Corrective Action::
Selects the ports with the specified corrective action.

``[-cable-length <text>]`` - Cable Length::
Selects the ports with the specified cable length.

``[-cable-identifier <text>]`` - Cable Identifier::
Selects the ports with the specified cable identifier.

``[-cable-end-id {end_0|end_1}]`` - Cable End Identifier::
Selects the ports with the specified cable end identifier.

``[-connector-type {QSFP|QSFP+|QSFP28|Mini-SAS HD|SFP}]`` - Connector
Type::
Selects the ports with the specified connector type.

``[-connector-vendor <text>]`` - Connector Vendor::
Selects the ports with the specified connector vendor.
```

```
``[-connector-part-number <text>]`` - Connector Part Number::
Selects the ports with the specified connector part number.

``[-connector-serial-number <text>]`` - Connector Serial Number::
Selects the ports with the specified connector serial number.

``[-is-redundant {true|false}]`` - Is Redundant?::
Selects the ports that are redundant (true) or not redundant (false). A
port is redundant if all attached devices have a second path through
another port.

``[-in-use {true|false}]`` - In Use?::
Selects the ports that are in use (true) or not in use (false). A port is
in use if any devices are connected using this port.

== Examples

The following example displays information about all storage ports in the
cluster:
```

```
cluster1::> storage port show
Speed VLAN
Node Port Type Mode (Gb/s) State Status ID
-------------- ---- ----- ------- ------ -------- --------- ----
cluster1-01
4a SAS - 0 enabled offline -
4b SAS - 0 enabled offline -
4c SAS - 0 enabled offline -
4d SAS - 0 enabled offline -
e3a ENET network - - - -
e3b ENET storage 100 enabled online 30
e5a ENET storage 100 enabled online 30
e5b ENET network - - - -
cluster1-02
4a SAS - 0 enabled offline -
4b SAS - 0 enabled offline -
4c SAS - 0 enabled offline -
4d SAS - 0 enabled offline -
e3a ENET network - - - -
e3b ENET storage 100 enabled online 30
e5a ENET storage 100 enabled online 30
e5b ENET network - - - -
16 entries were displayed.
```

```
The following example displays detailed information about port e5a on node
node1:
```

```
cluster1::> storage port show -node cluster1-01 -port e5a
```

Node: cluster1-01
Port: e5a
Port Type: ENET
Mode: storage
Description: 40G/100G Ethernet Controller CX5
Firmware Revision: 16.23.1020
MAC Address: ec:0d:9a:65:e4:44
Is Dedicated: false
Serial Number: MT1730X00227
Connector Vendor: Molex Inc.
Connector Part Number: 112-00322
Connector Serial Number: 532120266
Port Speed: 100 Gb/s
Port State: enabled
Port Status: online

```
:leveloffset: -1


= storage raid-options commands

:leveloffset: +1


[[IDc79675a44eaf8f0edbcff06cfb489cc7]]
= storage raid-options modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a RAID option

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage raid-options modify` command is used to modify the available
RAID options for each node in a cluster. The options are described in the
`storage raid-options` manual page.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node on which the RAID option is to be
modified.
```

```
``-name <text>`` - Option Name::
This parameter specifies the RAID option to be modified. To see the list
of RAID options that can be modified, use the xref:{relative_path}storage-
raid-options-show.html[storage raid-options show] command.

``[-value <text>]`` - Option Value::
This parameter specifies the value of the selected RAID option.

== Related Links

* xref:{relative_path}storage-raid-options-show.html[storage raid-options
show]




[[ID02edcafda1b7e48095b6484e19077179]]
= storage raid-options show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display RAID options

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage raid-options show` command displays information about all the
RAID options in a cluster. The options are described in the `storage raid-
options` manual page.

== 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>|local}]`` - Node::
Selects information about all the RAID options on the specified node.

``[-name <text>]`` - Option Name::
Selects information about the RAID options that have the specified name.

``[-value <text>]`` - Option Value::
Selects information about all the RAID options that have the specified
value.

``[-recommended-value <text>]`` - Recommended Value::
Selects information about all the RAID options that have the specified
recommended value.

``[-constraint <text>]`` - Option Constraint::
Selects information about all the RAID options that have the specified
constraint. The 'constraint' field indicates the expected setting for a
RAID option across both nodes of an HA pair. The possible values are:
+
  * ``_none_`` - no constraint on the value of this RAID option; nodes can
have different values
  * ``_same_preferred_`` - the same value should be used on both nodes of
an HA pair, otherwise the next takeover may not function correctly
  * ``_same_required_`` - the same value must be used on both nodes of an
HA pair, otherwise the next takeover will not function correctly
  * ``_only_one_`` - the same value should be used on both nodes of an HA
pair. If the values are different and a takeover is in progress, the value
of the RAID option on the node that is taking over will be used
  * ``_unknown_`` - no information about constraints for this RAID option



== Examples

The following example shows the RAID scrub settings for a node named
node1:
```

cluster1::> storage raid-options show -node node1 -name raid.scrub*
Current Recommended
Node Option Value Value Constraint
-------- ---------- ---------- ---------- ----------
node1 raid.scrub.perf_impact low low only_one
node1 raid.scrub.enable on off none
2 entries were displayed.

```
:leveloffset: -1
```

```
= storage raidlm commands

:leveloffset: +1


[[ID51bb46777550fd6b7aa1ad665ccc5bd6]]
= storage raidlm policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable/Disable the policy

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage raidlm policy modify` command is used to modify the available
Raid Layout Manager policy.

== Parameters

``-node <nodename>`` - Node names [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the node on which this Raid Layout Manager policy
is enabled.

``-policy-name <text>`` - Policy name for Raid-lm [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the Raid Layout Manager policy name.

``-policy-type <text>`` - Policy type for Raid-lm [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the Raid Layout Manager policy type.

``[-is-enabled {true|false}]`` - Is Policy Enabled? [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the value of the Raid Layout Manager policy to be
modified. To see the list of Raid Layout Manager policy that can be
modified, use the xref:{relative_path}storage-raidlm-policy-
show.html[storage raidlm policy show] command.

== Related Links
```

```
* xref:{relative_path}storage-raidlm-policy-show.html[storage raidlm
policy show]




[[IDf1b8b7307cac0f4f96bb168a343d98a8]]
= storage raidlm policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the policies

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage raidlm policy show` command displays information about Raid
Layout Manager policies and its value. Raid Layout Manager is a user space
process which polls system configuration, detect conflict if it deviates
from NetApp recommended best practices, and resolve conflict.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - Node names [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node on which this Raid Layout Manager policy
is enabled.

``[-policy-name <text>]`` - Policy name for Raid-lm [.advaced-
priv]#(privilege: advanced)#::
All the Raid Layout Manager policy.

``[-policy-type <text>]`` - Policy type for Raid-lm [.advaced-
```

```
priv]#(privilege: advanced)#::
Policy type of the Raid Layout Manager policy.

``[-is-enabled {true|false}]`` - Is Policy Enabled? [.advaced-
priv]#(privilege: advanced)#::
Value of the Raid Layout Manager policy that have the specified name.

== Examples

The following example shows the RAID-LM policy settings for node1 and
node2:
```

cluster1::> storage raidlm policy show -node node* -policy-name auto_unpartition_on_spares_low
Node Policy Type Policy Name Is Enabled
--- --- --- ---
node1
Shared-Disk auto_unpartition_on_spares_low true
node2
Shared-Disk auto_unpartition_on_spares_low true
2 entries were displayed.

```
:leveloffset: -1


= storage shelf commands

:leveloffset: +1


[[ID6bfee09d5fbebbef6c5c3f8069af750]]
= storage shelf show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of storage shelves

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage shelf show` command displays information about all the
storage shelves in the storage system. If no parameters are specified, the
default command displays the following information about the storage
```

```
shelves:

* Shelf Name
* Shelf ID
* Serial Number
* Model
* Module Type
* Status

To display detailed profile information about a single storage shelf, use
the `-shelf` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all the storage shelves, in column style
output.

| ``[-bay  ]``::
Displays the following details about the disk bays in the storage shelf:
+
  * The unique positional identifier of the disk bay
  * Whether a disk drive is installed in the bay
  * Bay type
  * Operational status of the disk bay



| ``[-connectivity  ]``::
Displays the following details about the connectivity from the node to the
storage shelf:
+
  * Node name
  * Initiator side switch port
  * Target side switch port
  * World-wide port name
  * Target Port Group Number (TPGN)



| ``[-cooling  ]``::
Displays the following details about the cooling elements and temperature
sensors of the storage shelf:
+
  * Element ID of the cooling fan
  * The current speed of the cooling fan in revolutions per minute (rpm)
  * Operational status of the cooling fan
```

```
  * Sensor ID of the temperature sensor element
  * Temperature at the sensor in degrees Celsius
  * Whether the current temperature at the sensor is the ambient
temperature
  * Low critical threshold value for the temperature sensor
  * Low warning threshold value for the temperature sensor
  * High critical threshold value for the temperature sensor
  * High warning threshold value for the temperature sensor
  * Operational status for the temperature sensor



| ``[-errors  ]``::
Displays the following error status information about the storage shelves
that have errors:
+
  * Error type
  * Error description



| ``[-module  ]``::
Displays the following details about the I/O modules attached to the
storage shelf:
+
  * Module ID
  * Module part number
  * Serial number of the Enclosure Services Controller Electronics element
  * Whether monitoring is enabled on this module
  * Whether this module is the SAS expander master module
  * Whether this module is the element reporting
  * Version of the firmware installed on the module
  * Latest firmware revision
  * Number of times, since the last boot, that this module has been
swapped
  * Operational status of the module



| ``[-port  ]``::
Displays the following details about the storage shelf ports:
+
  * Expander phy element identifier
  * SAS shelf port type
  * World-wide Port Name of the SAS port
  * Operational physical link rate of the SAS port in Gb/s
```

```
  * Negotiated physical link rate of the SAS port in Gb/s
  * Power status of the SAS port
  * Status of the SAS port
  * Fibre Channel shelf port ID
  * Fibre Channel shelf port type
  * Fibre Channel shelf port status



| ``[-power  ]``::
Displays the following details about the power supplies, voltage sensors,
and current sensors of the storage shelf:
+
  * Power Supply Unit (PSU) number
  * PSU type
  * PSU part number
  * PSU serial number
  * PSU power rating in watts
  * PSU crest factor
  * Power drawn from the PSU in watts
  * Whether the PSU can be reset via software control
  * Whether the auto power reset of the PSU is enabled
  * PSU firmware revision
  * Operational status of the PSU
  * Voltage sensor number
  * Voltage detected by the voltage sensor, in volts (V)
  * Operational status of the voltage sensor
  * Current sensor number
  * Current detected by the current sensor, in milliamps (mA)
  * Operational status of the current sensor



| ``[-instance  ]``::
Displays expanded information about all the storage shelves in the system.

``[-shelf <text>]`` - Shelf Name::
Displays information only about the storage shelves that match the names
you specify.

``[-node {<nodename>|local}]`` - Node::
Displays information only about the storage shelves that are attached to
the node you specify.

``[-shelf-uid <text>]`` - Shelf UID::
Displays information only about the storage shelf that matches the shelf
```

```
UID you specify. Example: `50:05:0c:c0:02:10:64:26 `

``[-stack-id {<integer>|-}]`` - Stack ID::
Displays information only about the storage shelves that are attached to
the stack that matches the stack ID you specify

``[-shelf-id <text>]`` - Shelf ID::
Displays information only about the storage shelves that match the shelf
ID you specify.

``[-module-type
{unknown|atfcx|esh4|iom3|iom6|iom6e|iom12|iom12e|iom12f|nsm100|nsm8e|psm3e
|iom12b|iom12g|nsm16e|iom12c|nsm100b}]`` - Shelf Module Type::
Displays information only about the storage shelves that match the module-
type you specify.

``[-connection-type {unknown|fc|sas|nvme}]`` - Shelf Connection Type::
Displays information only about the storage shelves that match the
connection type you specify. Example: FC or SAS.

``[-is-local-attach {true|false}]`` - Is the Shelf Local to This
Cluster?::
Displays information only about the storage shelves that are local (TRUE)
or remote (FALSE) to this cluster.

``[-vendor <text>]`` - Shelf Vendor::
Displays information only about the storage shelves that match the vendor
you specify.

``[-product-id <text>]`` - Shelf Product Identification::
Displays information only about the storage shelves that match the product
ID you specify.

``[-serial-number <text>]`` - Shelf Serial Number::
Displays information only about the storage shelf that matches the serial
number you specify.

``[-disk-count {<integer>|-}]`` - Disk Count::
Displays information only about the storage shelves that have the disk
count you specify.

``[-state {unknown|no-status|init-required|online|offline|missing}]`` -
Shelf State::
Displays information only about the storage shelves that are in the state
you specify.

``[-op-status {unknown|normal|warning|error|critical|standby-power}]`` -
```

```
Shelf Operational Status::
Displays information only about the storage shelves that are currently
operating under the status condition you specify.

``[-bay-id {<integer>|-}]`` - Bay ID::
Displays information only about the storage shelves that have bays that
match the bay ID you specify.

``[-bay-type {unknown|single-disk|multi-lun}]`` - Bay Type::
Displays information only about the storage shelves that have bays that
match the type of bay you specify.

``[-bay-has-disk {true|false}]`` - Bay Has Disk::
Displays information only about the storage shelves that have bays with
disk drives inserted in them (true) or empty bays (false).

``[-bay-op-status {unknown|normal|error}]`` - Bay Operational Status::
Displays information only about the storage shelves that have bays that
match the operational state you specify.

``[-controller {<nodename>|local}]`` - Controller Name::
Displays information only about the storage shelves that are connected to
the node you specify.

``[-controller-uuid <text>,...]`` - Controller UUID::
Displays information only about the storage shelves that are connected to
the node UUID you specify.

``[-initiator <text>,...]`` - Initiator::
Displays information only about the storage shelves that are visible to
the initiator you specify.

``[-initiator-wwpn <text>,...]`` - Initiator WWPN::
Displays information only about the storage shelves that are visible to
the initiator WWPN you specify.

``[-initiator-side-switch-port <text>,...]`` - Initiator Side Switch
Port::
Displays information only about the storage shelves that are visible to an
initiator connected to the switch port you specify.

``[-target-side-switch-port <text>,...]`` - Target Side Switch Port::
Displays information only about the storage shelves visible on target
ports identified by the switch port to which they are connected.

``[-target-port <text>,...]`` - Target Port::
Displays information only about the storage shelves visible on the
```

```
specified target ports identified by their World-Wide Port Name (WWPN).

``[-tpgn {<integer>|-}]`` - Target Port Group Number::
Displays information only about the storage shelves that belong to the
Target Port Group Name (TPGN) you specify.

``[-port-speed {<integer>|-}]`` - Port Speed::
Displays information only about the storage shelves with ports that match
the port speed you specify.

``[-io-kbps {<integer>|-}]`` - Kbytes/sec on Storage Shelf::
Displays information only about the storage shelves visible to an
initiator that has executed I/O at the throughput you specify.

``[-iops {<integer>|-}]`` - Number IOPS per Second on Storage Shelf::
Displays information only about the storage shelves visible to an
initiator that has executed the number of IOPs you specify.

``[-current-sensor-id {<integer>|-}]`` - Current Sensor ID::
Displays information only about the storage shelves with current sensor
that matches the current sensor ID you specify.

``[-current-sensor-location <text>,...]`` - Current Sensor Location::
Displays information only about the storage shelves with current sensors
installed at the location you specify.

``[-current-sensor-reading {<integer>|-}]`` - Current Sensor Reading::
Displays information only about the storage shelves with current sensors
that match the current reading you specify.

``[-current-op-status {unknown|normal|over-current-critical|under-current-
critical|not-supported|not-installed}]`` - Operational Status::
Displays information only about the storage shelves with current sensors
that match the operational status you specify.

``[-fan-id {<integer>|-}]`` - Fan ID::
Displays information only about the storage shelves with cooling fans that
match the fan IDs you specify.

``[-fan-location <text>,...]`` - Fan Location::
Displays information only about the storage shelves with cooling fans
installed.

``[-fan-rpm {<integer>|-}]`` - Fan Rotation Per Minute::
Displays information only about the storage shelves with cooling fans that
match the rpm rate you specify.
```

```
``[-fan-op-status {unknown|normal|off|error|not-supported|not-
installed}]`` - Fan Operational Status::
Displays information only about the storage shelves with cooling fans that
match the operational status you specify.

``[-module-id <text>,...]`` - Module ID::
Displays information only about the storage shelves with an I/O module
that matches the module ID you specify.

``[-module-location <text>,...]`` - Module Location::
Displays information only about the storage shelves with I/O modules in
the specified shelf module slots.

``[-module-part-number <text>,...]`` - Module Part Number::
Displays information only about the storage shelves with I/O modules that
match the module part numbers you specify.

``[-is-sas-master-module {true|false}]`` - Is SAS Expander Master
Module?::
Displays information only about the storage shelves with a SAS master I/O
module (true) or an I/O module that is not a SAS master (false). This
parameter applies only to SAS shelves.

``[-is-monitor-active {true|false}]`` - Is Monitor Active?::
Displays information only about the storage shelves whose monitoring is
enabled (true) or disabled (false).

``[-enclosure-type <text>,...]`` - Module Enclosure Type::
Displays information only about the storage shelves that match the
enclosure types you specify.

``[-es-serial-number <text>,...]`` - ES Electronics Element Serial
Number::
Displays information only about the storage shelves with I/O modules that
match the electronics serial numbers you specify.

``[-module-fru-id <text>,...]`` - Field Replaceable Unit ID::
Displays information only about the storage shelves with I/O modules that
match the field replaceable unit (FRU) IDs you specify.

``[-module-is-reporting-element {true|false}]`` - Is Reporting Element?::
Displays information only about the storage shelves with element reporting
I/O modules (true) or not (false).

``[-module-fw-revision <text>,...]`` - Firmware Revision::
Displays information only about the storage shelves with I/O modules that
match the firmware revision you specify.
```

```
``[-module-latest-fw-revision <text>,...]`` - Latest Firmware Revision::
Displays information only about the storage shelves with I/O modules that
match the latest firmware revision you specify.

``[-module-fw-progress {not-available|ready|in-progress|failed}]`` -
Module Firmware Progress::
Displays information only about the storage shelves with I/O modules that
match the specified firmware update progress.

``[-module-swap-count {<integer>|-}]`` - Module Swap Count::
Displays information only about the storage shelves whose I/O modules have
been swapped the specified number of times.

``[-module-op-status {unknown|normal|warning|error|not-installed}]`` -
Module Operational Status::
Displays information only about the storage shelves with I/O modules that
match the operational status you specify.

``[-sas-port-id <text>,...]`` - Port ID::
Displays information only about the storage shelves with SAS Ports that
match the port IDs you specify.

``[-sas-port-type
{unknown|circle|square|sil|disk|in|out|unused|host|dcm|aux1|aux2|hi_ho|a_t
o_b|b_to_a}]`` - Port Type::
Displays information only about the storage shelves with SAS Ports that
match the SAS port type you specify.

``[-sas-port-wwpn <text>,...]`` - Port World Wide Port Name::
Displays information only about the storage shelves with SAS Ports that
match the World-Wide Port Names you specify.

``[-sas-port-speed <text>,...]`` - Port Speed::
Displays information only about the storage shelves with SAS Ports that
match the port speed you specify.

``[-sas-negotiated-port-speed <text>,...]`` - Negotiated Port Speed::
Displays information only about the storage shelves with SAS Ports that
match the negotiated port speed you specify.

``[-sas-port-power-status <text>,...]`` - Port Power Status::
Displays information only about the storage shelves with SAS Ports that
match the power status you specify.

``[-sas-port-op-status {error|normal|off|unknown|byp-bad-term|bad-zone-
recovery|byp_clk_thr|byp_comma_los|byp_crc_brst_thr|byp_data_timeout|byp_d
```

```
rv_fault|byp_drv_pcycle|byp_drv_pwr|byp_drv_self|byp_gen|byp_init|byp_lip_
brst_thr|byp_lip_f8|byp_lip_rate_thr|byp_lipf7|byp_ltbi|byp_man|byp_no_dri
ve|byp_osc|byp_other_thr|byp_rec_los|byp_rport|byp_stall_thr|byp_wrd_brst_
thr|byp_wrd_rate_thr|byp_xmit_fault|diag_transmit|inserted|loopback|status
_unknown|warn_high_clk_delta|warn_high_crc_rate|warn_high_lip|warn_high_wr
d_rate|term|phy_dis_clk_fault|phy_dis_crc_err|phy_dis_crc_err_burst|phy_di
s_disparity|phy_dis_disparity_burst|phy_dis_emulate_reserve|phy_dis_inval_
dword|phy_dis_inval_dword_burst|phy_dis_loss_dword|phy_dis_loss_dword_burs
t|phy_dis_man_smp|phy_dis_manual|phy_dis_mirrored|empty|phy_dis_phy_change
|phy_dis_phy_change_burst|phy_dis_phy_reset|phy_dis_phy_reset_burst|phy_di
s_phy_unused|phy_ena|phy_ena_not_attach|phy_ena_unknown|phy_unknown|phy_di
s_illegal}]`` - Port Operational Status::
Displays information only about the storage shelves with SAS Ports that
match the operational status you specify.

``[-sas-port-module-id {A|B}]`` - Port Module ID::
Displays information only about the storage shelves with SAS Ports that
match the module ID you specify.

``[-fc-port-id <text>,...]`` - Fibre Channel Port ID::
Displays information only about the storage shelves with FC Ports that
match the port IDs you specify.

``[-fc-port-mode
{unknown|circle|square|sil|disk|in|out|unused|host|dcm|aux1|aux2|hi_ho|a_t
o_b|b_to_a}]`` - Fibre Channel Port Mode::
Displays information only about the storage shelves with FC Ports that
match the port modes you specify.

``[-fc-port-op-status {error|normal|off|unknown|byp-bad-term|bad-zone-
recovery|byp_clk_thr|byp_comma_los|byp_crc_brst_thr|byp_data_timeout|byp_d
rv_fault|byp_drv_pcycle|byp_drv_pwr|byp_drv_self|byp_gen|byp_init|byp_lip_
brst_thr|byp_lip_f8|byp_lip_rate_thr|byp_lipf7|byp_ltbi|byp_man|byp_no_dri
ve|byp_osc|byp_other_thr|byp_rec_los|byp_rport|byp_stall_thr|byp_wrd_brst_
thr|byp_wrd_rate_thr|byp_xmit_fault|diag_transmit|inserted|loopback|status
_unknown|warn_high_clk_delta|warn_high_crc_rate|warn_high_lip|warn_high_wr
d_rate|term|phy_dis_clk_fault|phy_dis_crc_err|phy_dis_crc_err_burst|phy_di
s_disparity|phy_dis_disparity_burst|phy_dis_emulate_reserve|phy_dis_inval_
dword|phy_dis_inval_dword_burst|phy_dis_loss_dword|phy_dis_loss_dword_burs
t|phy_dis_man_smp|phy_dis_manual|phy_dis_mirrored|empty|phy_dis_phy_change
|phy_dis_phy_change_burst|phy_dis_phy_reset|phy_dis_phy_reset_burst|phy_di
s_phy_unused|phy_ena|phy_ena_not_attach|phy_ena_unknown|phy_unknown|phy_di
s_illegal}]`` - Fibre Channel Port Operational Status::
Displays information only about the storage shelves with FC Ports that
match the operational status you specify.
```

```
``[-psu-id {<integer>|-}]`` - Power Supply Unit ID::
Displays information only about the storage shelves with power supply
units (PSU) that match the unit IDs you specify.

``[-psu-location <text>,...]`` - Power Supply Unit Location::
Displays information only about the storage shelves with PSUs that are
located at the specified location inside the shelf.

``[-psu-type <text>,...]`` - Power Supply Unit Type::
Displays information only about the storage shelves with PSUs that match
the PSU types you specify.

``[-psu-part-number <text>,...]`` - Power Supply Unit Part Number::
Displays information only about the storage shelves with PSUs that match
the PSU part number you specify.

``[-psu-serial-number <text>,...]`` - Power Supply Unit Serial Number::
Displays information only about the storage shelves with PSUs that match
the PSU serial numbers you specify.

``[-psu-reset-capable {true|false}]`` - Power Supply Unit Reset
Capability::
Displays information only about the storage shelves with reset capable
PSUs (true) or reset incapable PSUs (false).

``[-psu-is-enabled {true|false}]`` - Power Supply Unit Enable/Disable
Status::
Displays information only about the storage shelves with PSUs that are
enabled (true) or disabled (false).

``[-psu-fw-version <text>,...]`` - Power Supply Unit Firmware Version::
Displays information only about the storage shelves with PSUs that have
the firmware version you specify.

``[-psu-op-status {unknown|normal|error|dc-over-voltage|dc-under-
voltage|dc-over-current|over-temperature-error|failed|off|not-
supported|not-installed}]`` - Operational Status::
Displays information only about the storage shelves with PSUs that match
the operational status you specify.

``[-psu-power-rating {<integer>|-}]`` - Power Supply Power Ratings In
Watts::
Displays information only about the storage shelves with PSUs that match
the power rating you specify.

``[-psu-crest-factor {<integer>|-}]`` - Power Supply Crest Factor::
Displays information only about the storage shelves with PSUs that match
```

```
the crest factor value you specify.

``[-psu-power-drawn {<integer>|-}]`` - Power Drawn From PSU In Watts::
Displays information only about the storage shelves with PSUs that match
the drawn power you specify.

``[-temp-sensor-id {<integer>|-}]`` - Sensor Name::
Displays information only about the storage shelves with temperature
sensors that match the sensor IDs you specify.

``[-temp-sensor-location <text>,...]`` - Sensor Location::
Displays information only about the storage shelves with temperature
sensors that match the specified sensor locations inside the shelf.

``[-temp-sensor-reading {<integer>|-}]`` - Temperature Reading::
Displays information only about the storage shelves with temperature
sensors that match the temperature reading you specify.

``[-temp-is-ambient {true|false}]`` - Temperature Reading at Ambient
Value::
Displays information only about the storage shelves with temperature
sensors whose current temperature reading is ambient (true) or not
(false).

``[-temp-high-critical-threshold {<integer>|-}]`` - High Critical
Threshold::
Displays information only about the storage shelves with temperature
sensors that match the high critical threshold you specify.

``[-temp-high-warning-threshold {<integer>|-}]`` - High Warning
Threshold::
Displays information only about the storage shelves with temperature
sensors that match the high warning threshold you specify.

``[-temp-low-warning-threshold {<integer>|-}]`` - Low Warning Threshold::
Displays information only about the storage shelves with temperature
sensors that match the low warning threshold you specify.

``[-temp-low-critical-threshold {<integer>|-}]`` - Low Critical
Threshold::
Displays information only about the storage shelves with temperature
sensors that match the low critical threshold you specify.

``[-temp-op-status {unknown|normal|under-temperature|over-
temperature|error|not-supported|not-installed}]`` - Operational Status::
Displays information only about the storage shelves with temperature
sensors that match the operational status you specify.
```

```
``[-voltage-sensor-id {<integer>|-}]`` - Voltage Sensor ID::
Displays information only about the storage shelves with voltage sens- ors
that match the sensor IDs you specify.

``[-voltage-sensor-location <text>,...]`` - Voltage Sensor Location::
Displays information only about the storage shelves with voltage sens- ors
that match the specified sensor locations inside the shelf.

``[-voltage-sensor-reading <text>,...]`` - Voltage Current Reading::
Displays information only about the storage shelves with voltage sens- ors
that match the voltage reading you specify.

``[-voltage-op-status {unknown|normal|over-voltage-critical|under-voltage-
critical|not-supported|not-installed|not-recoverable}]`` - Operational
Status::
Displays information only about the storage shelves with voltage sens- ors
that match the operational status you specify.

``[-nsm-port-module-id {A|B}]`` - Port Module ID::
Displays information only about the storage shelves with PCIe Ports from
the specified module.

``[-nsm-port-id <integer>,...]`` - Port ID::
Displays information only about the storage shelves with PCIe Ports that
match the specified ID.

``[-nsm-port-type {cpu|disk|cx5|ethernet}]`` - Port Type::
Displays information only about the storage shelves with PCIe Ports that
match the specified type.

``[-nsm-port-state {ok|off-link-disabled|off-dll-link|link-down|no-
drive}]`` - Port State::
Displays information only about the storage shelves with PCIe Ports that
match the specified state.

``[-nsm-port-bay <integer>,...]`` - Port Bay::
Displays information only about the storage shelves with PCIe Ports that
match the specified bay.

``[-nsm-port-disk-id <integer>,...]`` - Port Disk ID::
Displays information only about the storage shelves with PCIe Ports that
match the specified disk ID.

``[-nsm-port-is-installed {true|false}]`` - Port Is Disk Installed::
Displays information only about the storage shelves with PCIe Ports that
have a disk installed.
```

```
``[-nsm-port-is-error {true|false}]`` - Port Has Error::
Displays information only about the storage shelves with PCIe Ports that
have errors.

``[-nsm-port-speed {2.5|5.0|8.0|16.0|32.0}]`` - Port Speed::
Displays information only about the storage shelves with PCIe Ports that
match the specified speed.

``[-nsm-port-speed-max {2.5|5.0|8.0|16.0|32.0}]`` - Max Port Speed::
Displays information only about the storage shelves with PCIe Ports that
match the specified maximum speed.

``[-nsm-port-lane-width <integer>,...]`` - Port Lane Width::
Displays information only about the storage shelves with PCIe Ports that
match the specified lane width.

``[-nsm-port-lane-width-max <integer>,...]`` - Max Port Lane Width::
Displays information only about the storage shelves with PCIe Ports that
match the specified maximum lane width.

``[-dimm-module-id {A|B}]`` - DIMM Module ID::
Displays information only about the storage shelves with DIMMs from the
specified module.

``[-dimm-id <integer>,...]`` - DIMM ID::
Displays information only about the storage shelves with DIMMs that match
the specified ID.

``[-dimm-serial-number <text>,...]`` - DIMM Serial Number::
Displays information only about the storage shelves with DIMMs that match
the specified serial number.

``[-dimm-part-number <text>,...]`` - DIMM Part Number::
Displays information only about the storage shelves with DIMMs that match
the specified part number.

``[-dimm-vendor <text>,...]`` - DIMM Vendor::
Displays information only about the storage shelves with DIMMs that match
the specified vendor.

``[-dimm-type <text>,...]`` - DIMM Type::
Displays information only about the storage shelves with DIMMs that match
the specified type.

``[-dimm-size <text>,...]`` - DIMM Size::
Displays information only about the storage shelves with DIMMs that match
```

```
the specified size.

``[-dimm-speed <text>,...]`` - DIMM Speed::
Displays information only about the storage shelves with DIMMs that match
the specified speed.

``[-dimm-location <text>,...]`` - DIMM Location::
Displays information only about the storage shelves with DIMMs that match
the specified location.

``[-dimm-op-status {unknown|normal|error|not-supported|not-installed}]`` -
DIMM Operational Status::
Displays information only about the storage shelves with DIMMs that match
the specified operational status.

``[-boot-device-module-id {A|B}]`` - Boot Device Module ID::
Displays information only about the storage shelves with boot devices from
the specified module.

``[-boot-device-id <integer>,...]`` - Boot Device ID::
Displays information only about the storage shelves with boot devices that
match the specified ID.

``[-boot-device-serial-number <text>,...]`` - Boot Device Serial Number::
Displays information only about the storage shelves with boot devices that
match the specified serial number.

``[-boot-device-part-number <text>,...]`` - Boot Device Part Number::
Displays information only about the storage shelves with boot devices that
match the specified part number.

``[-boot-device-vendor <text>,...]`` - Boot Device Vendor::
Displays information only about the storage shelves with boot devices that
match the specified vendor.

``[-boot-device-type <text>,...]`` - Boot Device Type::
Displays information only about the storage shelves with boot devices that
match the specified type.

``[-boot-device-size <text>,...]`` - Boot Device Size::
Displays information only about the storage shelves with boot devices that
match the specified size.

``[-boot-device-op-status {unknown|normal|error|not-supported|not-
installed}]`` - Boot Device Operational Status::
Displays information only about the storage shelves with boot devices that
match the specified operational status.
```

```
``[-coin-battery-module-id {A|B}]`` - Coin Battery Module ID::
Displays information only about the storage shelves with coin batteries
from the specified module.

``[-coin-battery-id <integer>,...]`` - Coin Battery ID::
Displays information only about the storage shelves with coin batteries
that match the specified ID.

``[-coin-battery-voltage <integer>,...]`` - Coin Battery Voltage (mV)::
Displays information only about the storage shelves with coin batteries
that match the specified voltage.

``[-coin-battery-op-status {unknown|normal|error|low|high|not-
supported|not-installed}]`` - Coin Battery Operational Status::
Displays information only about the storage shelves with coin batteries
that match the specified operational status.

``[-error-type
{Unknown|ACPP|Bootdevice|Coinbattery|Configuration|Current|Dimm|Expander|F
an|Module|PCM|Power|Temperature|Voltage}]`` - Error Type::
Displays information only about the storage shelves with errors that match
the error type you specify.

``[-error-severity {unknown|notice|warning|error|critical}]`` - Error
Severity::
Displays information only about the storage shelves with errors that match
the error severity you specify.

== Examples

The following example displays information about all storage shelves:
```

```
cluster1::> storage shelf show
Module Operational
Shelf Name Shelf ID Serial Number Model Type Status
---------------- -------- -------------- ---------- ------ -----------
1.1 1 6000832415 DS2246 IOM6 Critical
1.2 2 6000647652 DS2246 IOM6 Normal
1.3 3 6000003844 DS2246 IOM6 Normal
1.4 4 SHJ000000013A9E DS4246 IOM6 Normal
1.5 5 SHJ000000013A84 DS4246 IOM6 Normal
1.6 6 6000005555 DS2246 IOM6 Normal
6 entries were displayed.
cluster1::>
```

```
The following example displays expanded information about a storage shelf
named 1.2:
```

```
cluster1::> storage shelf show -shelf 1.2 -instance
Shelf Name: 1.2
Stack ID: 1
Shelf ID: 2
Shelf UID: 50:0a:09:80:01:b9:75:41
Serial Number: 6000647652
Module Type: IOM6
Model: DS2246
Shelf Vendor: NETAPP
Disk Count: 12
Connection Type: SAS
Shelf State: Online
Status: Normal
Modules: Module is
Monitor Is Reporting FW Update Latest Swap Operational Module
ID Part No. ES Serial No. is Active Master Element Progress FW Rev. FW Rev. Count Status Location
--- ------------ --------------- --------- ------ --------- ------------- ------- ------- ----- ----------- --------------
a 111-00190+A0 8006437891 true false false not-available 0191 - 0 normal rear of the shelf at the top left
b 111-00190+A0 8006435180 true true true not-available 0191 - 0 normal rear of the shelf at the top right
Paths:
Speed
Controller Initiator Initiator Side Switch Port Target Side Switch Port Target Port TPGN Gb/s I/O KB/s IOPS
------------------ --------- -------------------------- -------------------------- ------------------ ------ ----- -------- -------
stsw-8020-01 0a - - - - - - -
stsw-8020-01 2b - - - - - - -
stsw-8020-02 0a - - - - - - -
stsw-8020-02 2b - - - - - - -
Power Supply Units:
Crest Power Reset PSU Operational
ID Type Part# Serial# Power Rating Factor Drawn Capable Enabled Firmware Status PSU Location
--- ---- ------------ --------------- -------------- ------ ------- ------- ------- -------- ---------------- ------------------------
1 9C 114-00065+A1 XXT131052637 - - - false true 020F normal rear of the shelf at the bottom left
2 9C 114-00065+A1 XXT131052551 - - - false true 020F normal rear of the shelf at the bottom right
Voltage Sensors:
Voltage Operational
ID (V) Status Sensor Location
--- ------- ---------------------- ------------------------
1 5.70 normal rear of the shelf on the lower left power supply
2 12.300 normal rear of the shelf on the lower left power supply
3 5.70 normal rear of the shelf on the lower right power supply
4 12.180 normal rear of the shelf on the lower right power supply
Current Sensors:
Current Operational
ID (mA) Status Sensor Location
--- ------- ---------------------- ------------------------
1 0 normal rear of the shelf on the lower left power supply
2 0 normal rear of the shelf on the lower left power supply
3 0 normal rear of the shelf on the lower right power supply
4 0 normal rear of the shelf on the lower right power supply
```

Fans:
Speed Operational
ID (RPM) Status Fan Location
--- ------ ---------- ----------
1 3000 normal rear of the shelf on the lower left power supply
2 2970 normal rear of the shelf on the lower left power supply
3 3000 normal rear of the shelf on the lower right power supply
4 2970 normal rear of the shelf on the lower right power supply
Temperature:
— Thresholds °C — Temp Is Low Low High High Operational
ID °C Ambient Crit Warn Crit Warn Status Sensor Location
--- ---- ------- ---- ---- ---- ---- ---------- ----------
1 23 true 0 5 42 40 normal front of the shelf on the left, on the OPS panel
2 26 false 5 10 55 50 normal inside of the shelf on the midplane
3 24 false 5 10 55 50 normal rear of the shelf on the lower left power supply
4 39 false 5 10 70 65 normal rear of the shelf on the lower left power supply
5 25 false 5 10 55 50 normal rear of the shelf on the lower right power supply
6 36 false 5 10 70 65 normal rear of the shelf on the lower right power supply
7 25 false 5 10 60 55 normal rear of the shelf at the top left, on shelf module A
8 26 false 5 10 60 55 normal rear of the shelf at the top right, on shelf module B
SAS Ports:
— Port Speeds Gb/s — Power Port
Phy # IOM Port Type WWPN Operational Negotiated Status Status
----- --- --------- ---------- ---------- ---------- ------ ----------
0 A Square 500a098004b063b0 6.0 - - Enabled
1 A Square 500a098004b063b0 6.0 - - Enabled
2 A Square 500a098004b063b0 6.0 - - Enabled
3 A Square 500a098004b063b0 6.0 - - Enabled
4 A Circle 500a09800569f03f 6.0 - - Enabled
5 A Circle 500a09800569f03f 6.0 - - Enabled
6 A Circle 500a09800569f03f 6.0 - - Enabled
7 A Circle 500a09800569f03f 6.0 - - Enabled
8 A Disk 500605ba00c1cb8d 6.0 6.0 on Enabled
9 A Disk 500605ba00c1ea8d 6.0 6.0 on Enabled
10 A Disk 500605ba00c1d111 6.0 6.0 on Enabled
11 A Disk 500605ba00c1bc49 6.0 6.0 on Enabled
12 A Disk 500605ba00c1cdfd 6.0 6.0 on Enabled
13 A Disk 500605ba00c1c531 6.0 6.0 on Enabled
14 A Disk 500605ba00c1eb05 6.0 6.0 on Enabled
15 A Disk 500605ba00c1ec29 6.0 6.0 on Enabled
16 A Disk 500605ba00c1bc29 6.0 6.0 on Enabled
17 A Disk 500605ba00c1c471 6.0 6.0 on Enabled
18 A Disk 500605ba00c039a9 6.0 6.0 on Enabled
19 A Disk 500605ba00c1c4dd 6.0 6.0 on Enabled
20 A Disk - - - - Empty
21 A Disk - - - - Empty
22 A Disk - - - - Empty
23 A Disk - - - - Empty
24 A Disk - - - - Empty
25 A Disk - - - - Empty
26 A Disk - - - - Empty
27 A Disk - - - - Empty
28 A Disk - - - - Empty
29 A Disk - - - - Empty

30 A Disk - - - - Empty
31 A Disk - - - - Empty
32 A SIL - - - - Disabled
33 A SIL - - - - Disabled
34 A SIL - - - - Disabled
35 A SIL - - - - Disabled
0 B Square 500a098004af9e30 6.0 - - Enabled
1 B Square 500a098004af9e30 6.0 - - Enabled
2 B Square 500a098004af9e30 6.0 - - Enabled
3 B Square 500a098004af9e30 6.0 - - Enabled
4 B Circle 500a098005688dbf 6.0 - - Enabled
5 B Circle 500a098005688dbf 6.0 - - Enabled
6 B Circle 500a098005688dbf 6.0 - - Enabled
7 B Circle 500a098005688dbf 6.0 - - Enabled
8 B Disk 500605ba00c1cb8e 6.0 6.0 on Enabled
9 B Disk 500605ba00c1ea8e 6.0 6.0 on Enabled
10 B Disk 500605ba00c1d112 6.0 6.0 on Enabled
11 B Disk 500605ba00c1bc4a 6.0 6.0 on Enabled
12 B Disk 500605ba00c1cdfe 6.0 6.0 on Enabled
13 B Disk 500605ba00c1c532 6.0 6.0 on Enabled
14 B Disk 500605ba00c1eb06 6.0 6.0 on Enabled
15 B Disk 500605ba00c1ec2a 6.0 6.0 on Enabled
16 B Disk 500605ba00c1bc2a 6.0 6.0 on Enabled
17 B Disk 500605ba00c1c472 6.0 6.0 on Enabled
18 B Disk 500605ba00c039aa 6.0 6.0 on Enabled
19 B Disk 500605ba00c1c4de 6.0 6.0 on Enabled
20 B Disk - - - - Empty
21 B Disk - - - - Empty
22 B Disk - - - - Empty
23 B Disk - - - - Empty
24 B Disk - - - - Empty
25 B Disk - - - - Empty
26 B Disk - - - - Empty
27 B Disk - - - - Empty
28 B Disk - - - - Empty
29 B Disk - - - - Empty
30 B Disk - - - - Empty
31 B Disk - - - - Empty
32 B SIL - - - - Disabled
33 B SIL - - - - Disabled
34 B SIL - - - - Disabled
35 B SIL - - - - Disabled
FC Ports:
Port
ID Port Type Status
--- --------- ----------
- - -
Bays:
Has Operational
ID Disk Bay Type Status
--- ----- ----------- ----------
0 true single-disk normal
1 true single-disk normal
2 true single-disk normal

3 true single-disk normal
4 true single-disk normal
5 true single-disk normal
6 true single-disk normal
7 true single-disk normal
8 true single-disk normal
9 true single-disk normal
10 true single-disk normal
11 true single-disk normal
12 false single-disk normal
13 false single-disk normal
14 false single-disk normal
15 false single-disk normal
16 false single-disk normal
17 false single-disk normal
18 false single-disk normal
19 false single-disk normal
20 false single-disk normal
21 false single-disk normal
22 false single-disk normal
23 false single-disk normal
cluster1::>

```
The following example displays information about the power supplies,
voltage sensors and current sensors of the storage shelf 1.1:
```

cluster1::> storage shelf show -shelf 1.1 -power
Shelf Name: 1.1
Stack ID: 1
Shelf ID: 1
Shelf UID: 50:0a:09:80:01:cb:d6:84
Serial Number: 6000832415
Module Type: IOM6
Model: DS2246
Shelf Vendor: NETAPP
Disk Count: 12
Connection Type: SAS
Shelf State: Online
Status: Normal
Power Supply Units:
Crest Power Reset PSU Operational
ID Type Part# Serial# Power Rating Factor Drawn Capable Enabled Firmware Status
--- ---- ---------- --------- ---------- ------ ------ ------ ------ -------- ---------
1 9C 114-00065+A1 XXT132835072 - - - false true 020F normal
2 9C 114-00065+A1 XXT132835073 - - - false true 020F normal
Voltage Sensors:
Voltage Operational
ID (V) Status
--- ------ ---------
1 5.70 normal
2 12.180 normal
3 5.70 normal

```
4 12.300 normal
Current Sensors:
Current Operational
ID (mA) Status
--- ------- ----------------------
1 0 normal
2 0 normal
3 3900 normal
4 0 normal
Errors:
------
Critical condition is detected in storage shelf power supply unit "1". The unit might fail.
Critical over temperature failure for temperature sensor "1". Current temperature: "75" C ("167" F).
cluster1::>
```

```
The following example displays information about the cooling elements and
temperature sensors inside the storage shelf 1.2:
```

```
cluster1::> storage shelf show -shelf 1.2 -cooling
Shelf Name: 1.2
Stack ID: 1
Shelf ID: 2
Shelf UID: 50:0a:09:80:01:b9:75:41
Serial Number: 6000647652
Module Type: IOM6
Model: DS2246
Shelf Vendor: NETAPP
Disk Count: 12
Connection Type: SAS
Shelf State: Online
Status: Normal
Fans:
Speed Operational
ID (RPM) Status
— ----- --------------
1 3000 normal
2 3000 normal
3 3000 normal
4 2970 normal
Temperature:
— Thresholds °C —  Temp Is Low Low High High Operational
ID °C Ambient Crit Warn Crit Warn Status
--- ---- ------- ---- ---- ---- ---- ------------------
1 23 true 0 5 42 40 normal
2 26 false 5 10 55 50 normal
3 24 false 5 10 55 50 normal
4 39 false 5 10 70 65 normal
5 25 false 5 10 55 50 normal
6 36 false 5 10 70 65 normal
7 25 false 5 10 60 55 normal
8 27 false 5 10 60 55 normal
Errors:
```

-
cluster1::>

```
The following example displays information about the connectivity from the
node to the storage shelf 1.2:
```

cluster1::> storage shelf show -shelf 1.2 -connectivity
Shelf Name: 1.2
Stack ID: 1
Shelf ID: 2
Shelf UID: 50:0a:09:80:01:b9:75:41
Serial Number: 6000647652
Module Type: IOM6
Model: DS2246
Shelf Vendor: NETAPP
Disk Count: 12
Connection Type: SAS
Shelf State: Online
Status: Normal
Paths:
Controller Initiator Initiator Side Switch Port Target Side Switch Port Target Port TPGN
--- --- --- --- --- ---
stsw-8020-01 0a - - - -
stsw-8020-01 2b - - - -
stsw-8020-02 0a - - - -
stsw-8020-02 2b - - - -
Errors:

-
cluster1::>

```
The following example displays information about the disk bays of the
storage shelf 1.2:
```

cluster1::> storage shelf show -shelf 1.2 -bay
Shelf Name: 1.2
Stack ID: 1
Shelf ID: 2
Shelf UID: 50:0a:09:80:01:b9:75:41
Serial Number: 6000647652
Module Type: IOM6
Model: DS2246
Shelf Vendor: NETAPP
Disk Count: 12
Connection Type: SAS
Shelf State: Online
Status: Normal
Bays:
Has Operational

```
ID Disk Bay Type Status
--- ----- ----------- -----------
0 true single-disk normal
1 true single-disk normal
2 true single-disk normal
3 true single-disk normal
4 true single-disk normal
5 true single-disk normal
6 true single-disk normal
7 true single-disk normal
8 true single-disk normal
9 true single-disk normal
10 true single-disk normal
11 true single-disk normal
12 false single-disk normal
13 false single-disk normal
14 false single-disk normal
15 false single-disk normal
16 false single-disk normal
17 false single-disk normal
18 false single-disk normal
19 false single-disk normal
20 false single-disk normal
21 false single-disk normal
22 false single-disk normal
23 false single-disk normal
Errors:
------
-
cluster1::>
```

```
The following example displays information about the ports of the storage
shelf 1.2:
```

```
cluster1::> storage shelf show -shelf 1.2 -port
Shelf Name: 1.2
Stack ID: 1
Shelf ID: 2
Shelf UID: 50:0a:09:80:01:b9:75:41
Serial Number: 6000647652
Module Type: IOM6
Model: DS2246
Shelf Vendor: NETAPP
Disk Count: 12
Connection Type: SAS
Shelf State: Online
Status: Normal
SAS Ports:
 — Port Speeds Gb/s — Power Port
Phy # IOM Port Type WWPN Operational Negotiated Status Status
----- --- --------- ----------------- ----------- ---------- ------ -----------
0 A Square 500a098004b063b0 6.0 - - Enabled
```

1 A Square 500a098004b063b0 6.0 - - Enabled
2 A Square 500a098004b063b0 6.0 - - Enabled
3 A Square 500a098004b063b0 6.0 - - Enabled
4 A Circle 500a09800569f03f 6.0 - - Enabled
5 A Circle 500a09800569f03f 6.0 - - Enabled
6 A Circle 500a09800569f03f 6.0 - - Enabled
7 A Circle 500a09800569f03f 6.0 - - Enabled
8 A Disk 500605ba00c1cb8d 6.0 6.0 on Enabled
9 A Disk 500605ba00c1ea8d 6.0 6.0 on Enabled
10 A Disk 500605ba00c1d111 6.0 6.0 on Enabled
11 A Disk 500605ba00c1bc49 6.0 6.0 on Enabled
12 A Disk 500605ba00c1cdfd 6.0 6.0 on Enabled
13 A Disk 500605ba00c1c531 6.0 6.0 on Enabled
14 A Disk 500605ba00c1eb05 6.0 6.0 on Enabled
15 A Disk 500605ba00c1ec29 6.0 6.0 on Enabled
16 A Disk 500605ba00c1bc29 6.0 6.0 on Enabled
17 A Disk 500605ba00c1c471 6.0 6.0 on Enabled
18 A Disk 500605ba00c039a9 6.0 6.0 on Enabled
19 A Disk 500605ba00c1c4dd 6.0 6.0 on Enabled
20 A Disk - - - - Empty
21 A Disk - - - - Empty
22 A Disk - - - - Empty
23 A Disk - - - - Empty
24 A Disk - - - - Empty
25 A Disk - - - - Empty
26 A Disk - - - - Empty
27 A Disk - - - - Empty
28 A Disk - - - - Empty
29 A Disk - - - - Empty
30 A Disk - - - - Empty
31 A Disk - - - - Empty
32 A SIL - - - - Disabled
33 A SIL - - - - Disabled
34 A SIL - - - - Disabled
35 A SIL - - - - Disabled
0 B Square 500a098004af9e30 6.0 - - Enabled
1 B Square 500a098004af9e30 6.0 - - Enabled
2 B Square 500a098004af9e30 6.0 - - Enabled
3 B Square 500a098004af9e30 6.0 - - Enabled
4 B Circle 500a098005688dbf 6.0 - - Enabled
5 B Circle 500a098005688dbf 6.0 - - Enabled
6 B Circle 500a098005688dbf 6.0 - - Enabled
7 B Circle 500a098005688dbf 6.0 - - Enabled
8 B Disk 500605ba00c1cb8e 6.0 6.0 on Enabled
9 B Disk 500605ba00c1ea8e 6.0 6.0 on Enabled
10 B Disk 500605ba00c1d112 6.0 6.0 on Enabled
11 B Disk 500605ba00c1bc4a 6.0 6.0 on Enabled
12 B Disk 500605ba00c1cdfe 6.0 6.0 on Enabled
13 B Disk 500605ba00c1c532 6.0 6.0 on Enabled
14 B Disk 500605ba00c1eb06 6.0 6.0 on Enabled
15 B Disk 500605ba00c1ec2a 6.0 6.0 on Enabled
16 B Disk 500605ba00c1bc2a 6.0 6.0 on Enabled
17 B Disk 500605ba00c1c472 6.0 6.0 on Enabled
18 B Disk 500605ba00c039aa 6.0 6.0 on Enabled

19 B Disk 500605ba00c1c4de 6.0 6.0 on Enabled
20 B Disk - - - - Empty
21 B Disk - - - - Empty
22 B Disk - - - - Empty
23 B Disk - - - - Empty
24 B Disk - - - - Empty
25 B Disk - - - - Empty
26 B Disk - - - - Empty
27 B Disk - - - - Empty
28 B Disk - - - - Empty
29 B Disk - - - - Empty
30 B Disk - - - - Empty
31 B Disk - - - - Empty
32 B SIL - - - - Disabled
33 B SIL - - - - Disabled
34 B SIL - - - - Disabled
35 B SIL - - - - Disabled
FC Ports:
Port
ID Port Type Status
--- --------- -----------
- - -
Errors:

-
cluster1::>

```
The following example displays error information about the storage shelves
that have errors:
```

cluster1::> storage shelf show -errors
Shelf Name: 1.1
Shelf UID: 50:0a:09:80:01:cb:d6:84
Serial Number: 6000832415
Error Type Description
--- ---
Power Critical condition is detected in storage shelf power supply unit "1". The unit might fail.
Temperature Critical over temperature failure for temperature sensor "1". Current temperature: "75" C ("167" F).

```
[[ID4f6bec67fe02a2d306178c01e17eff42]]
= storage shelf acp configure
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Configure alternate control path (ACP)

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Configure the ACP connectivity on the cluster. Enabling ACP connectivity
is non-disruptive to the cluster.

== Parameters

``-is-enabled {true|false}`` - Is Enabled?::
Configures the connectivity to the specified state.

``[-subnet <IP Address>]`` - Subnet::
Configures the connectivity to the specified subnet.

``[-netmask <IP Address>]`` - Netmask::
Configures the connectivity to the specified netmask.

``[-channel {out-of-band|in-band}]`` - Channel::
Configures the connectivity to the specified channel.

== Examples

The following example configures out-of-band ACP connectivity on each
node:
```

cluster1::> storage shelf acp configure -is-enabled true -channel out-of-band -subnet 192.168.0.1 -netmask 255.255.255.0

```
The following example configures in-band ACP connectivity on each node:
```

cluster1::> storage shelf acp configure -is-enabled true -channel in-band

```
The following example disables ACP connectivity on each node:
```

cluster1::> storage shelf acp configure -is-enabled false

```
[[ID1e40ef5ffecaabb9a67ae29b27797b45]]
= storage shelf acp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show connectivity information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Displays information about the ACP connectivity on each node

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <field-name>", ... parameter, the command
output also includes the specified field or fields. You can use '-fields
?' to display the fields to specify.

| ``[-errors  ]``::
If you specify the -errors parameter, the command displays detailed
information about all modules with errors.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Selects the nodes that match this parameter value.

``[-is-enabled {true|false}]`` - Is Enabled?::
Selects the nodes that are enabled or disabled.

``[-port <text>]`` - Port::
Selects the nodes that match the specified port on which ACP is
configured.
```

```
``[-address <IP Address>]`` - IP Address::
Selects the nodes with the specified IP address.

``[-subnet <IP Address>]`` - Subnet::
Selects the nodes with the specified subnet.

``[-netmask <IP Address>]`` - Netmask::
Selects the nodes with the specified netmask.

``[-connection-status {no-connectivity|partial-connectivity|full-
connectivity|additional-connectivity|unknown-connectivity|not-
available|connection-disabled}]`` - Connection Status::
Selects the nodes with the specified connection status.

``[-error-id <integer>]`` - Error ID::
Selects the node with the specified error ID.

``[-error-type {No-Error|Connection-Issue|Connection-Activity|Module-
Error|Shelf-Error}]`` - Error Type::
The error type, in case of a connection error.

``[-error-severity {unknown|notice|warning|error|critical}]`` - Error
Severity::
The error severity, in case of a connection error.

``[-error-text <text>]`` - Error Text::
Selects the node with the specified error text.

``[-corrective-action <text>]`` - Corrective Action::
Selects the node with the specified corrective action.

``[-channel {unknown|out-of-band|in-band}]`` - Channel::
Selects the nodes that has channel configured out-of-band or in-band.

== Examples

The following example displays ACP connectivity on each node (in-band):
```

```
fas2750-2n-rtp-1::> storage shelf acp show
Node Channel Connectivity
----------------- ------------------- ---------------------
fas2750-rtp-1a in-band active
fas2750-rtp-1b in-band active
2 entries were displayed.
```

```
The following example displays ACP connectivity on each node (out of
band):
```

fas2750-2n-rtp-1::> storage shelf acp show
Node Channel Connectivity
--- --- ---
fas2750-rtp-1a out-of-band full-connectivity
fas2750-rtp-1b out-of-band full-connectivity
2 entries were displayed.

```
The following example displays the -instance output of the storage acp
show (in-band) command. Use this command to display details on
connectivity and configuration.
```

fas2750-2n-rtp-1::> storage shelf acp show -instance
Node: fas2750-rtp-1a
Channel: in-band
Enable Status: true
Connection Status: active
Node: fas2750-rtp-1b
Channel: in-band
Enable Status: true
Connection Status: active
2 entries were displayed.

```
The following example displays the -instance output of the storage acp
show (out-of-band) command. Use this command to display details on
connectivity and configuration.
```

fas2750-2n-rtp-1::> storage shelf acp show -instance
Node: fas2750-rtp-1a
Channel: out-of-band
Enable Status: true
Port: e0P
IP Address: 192.168.1.74
Subnet: 192.168.0.1
Netmask: 255.255.252.0
Connection Status: full-connectivity
Node: fas2750-rtp-1b
Channel: out-of-band
Enable Status: true
Port: e0P
IP Address: 192.168.1.75
Subnet: 192.168.0.1
Netmask: 255.255.252.0
Connection Status: full-connectivity
2 entries were displayed.

```
[[IDa4ed92f5c5b23ededa09a677bf24433f]]
= storage shelf acp module show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show modules connected to the cluster

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Displays information about the modules connected to each node

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <field-name>", ... parameter, the command
output also includes the specified field or fields. You can use '-fields
?' to dis play the fields to specify

| ``[-errors  ]``::
If you specify the -errors parameter, the command displays detailed
information about all modules with errors.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Selects the modules that match this parameter value.

``[-mac-address <text>]`` - MAC Address::
Selects the module that match the specified MAC address.

``[-module-name <text>]`` - Module name::
Selects the module that match the specified module name.

``[-module-address <IP Address>]`` - IP Address::
Selects the module that match the specified IP address.

``[-protocol-version <text>]`` - Protocol Version::
Selects the modules that match the specified protocol version.
```

```
``[-firmware-version <text>]`` - Firmware Version::
Selects the modules that match the specified firmware version.

``[-acpa-id <integer>]`` - ACPA assigner ID::
Selects the modules that match the specified ACPA ID.

``[-shelf-serial-number <text>]`` - Shelf Serial Number::
Selects the modules that match the specified shelf serial number.

``[-iom-type
{Unknown|iom3|iom6|iom6e|iom12|iom12e|iom12f|iom12b|iom12g|iom12c}]`` -
IOM Type::
Selects the modules that match the specified IOM type (IOM3/IOM6/IOM6E).

``[-last-contact <integer>]`` - Last Contact (secs)::
Selects the modules that match the specified last contact.

``[-state {unknown|initializing|discovery-complete|awaiting-inband|no-
inband|active|awaiting-bootp|updating-firmware|connection-error|firmware-
update-required|rebooting|fail|unsupported|degraded|shelf-off}]`` - Local
Node State::
Selects the modules that match the specified state.

``[-stack-id {<integer>|-}]`` - Stack ID::
Selects the modules that match the specified stack ID.

``[-shelf-id <text>]`` - Shelf ID::
Selects the modules that match the specified shelf ID.

``[-adapter-name <text>]`` - Adapter Name::
Selects the modules that match the specified adapter name.

``[-error-id <integer>,...]`` - Error ID::
Selects the modules that match the specified error ID.

``[-error-text <text>,...]`` - Error Text::
The error text, in case of a module error.

``[-corrective-action <text>,...]`` - Corrective Action::
The corrective action, in case of a module error.

``[-error-type {No-Error|Connection-Issue|Connection-Activity|Module-
Error|Shelf-Error}]`` - Error Type::
Selects the modules that match the specified error type.

``[-error-severity {unknown|notice|warning|error|critical}]`` - Error
Severity::
```

```
Selects the modules that match the specified error severity.

``[-power-cycle-count <integer>]`` - Power Cycle count::
Number of times a shelf power cycle has been performed on a shelf

``[-power-off-count <integer>]`` - Power Off count::
Number of times a shelf power off has been performed on a shelf

``[-power-on-count <integer>]`` - Power On count::
Number of times a shelf power on has been performed on a shelf

``[-expander-reset-count <integer>]`` - Expander reset count::
Number of times an expander reset has been performed on a module

``[-expander-power-cycle-count <integer>]`` - Expander power cycle count::
Number of times an expander power cycle has been performed on a module

== Examples

The following example displays the ACP modules connected to each node:
```

cluster1::> storage shelf acp module show
Node Module Name State
---------- ---------- ----------
stor-v4-1a-1b-01 1.10.A Active
1.10.B Active
1.254.B Active
1.254.A Active

stor-v4-1a-1b-02 1.10.A Active
1.10.B Active
1.254.B Active
1.254.A Active
8 entries were displayed.

```
The following example displays the -instance output of the storage shelf
acp module show. More details on each module can be seen here.
```

cluster1::> storage shelf acp module show -instance
Node: stor-v4-1a-1b-01
Module Name: 1.10.A
Mac Address: 00:a0:98:19:53:ee
IOM Type: IOM6E
Shelf Serial Number: SHJMS000000001A
IP Address: 192.168.3.239
Protocol Version: 2.1.1.21
Assigner ID: 2.1.1.21
State: Active

Last Contact: 203
Power Cycle Count: 0
Power Off Count: 0
Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-01
Module Name: 1.10.B
Mac Address: 00:a0:98:19:55:16
IOM Type: IOM6E
Shelf Serial Number: SHJMS000000001A
IP Address: 192.168.1.23
Protocol Version: 2.1.1.21
Assigner ID: 2.1.1.21
State: Active
Last Contact: 206
Power Cycle Count: 0
Power Off Count: 0
Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-01
Module Name: 1.254.B
Mac Address: 00:a0:98:32:d6:ac
IOM Type: IOM6
Shelf Serial Number: 6000368103
IP Address: 192.168.2.173
Protocol Version: 1.2.2. 8
Assigner ID: 1.2.2. 8
State: Active
Last Contact: 215
Power Cycle Count: 0
Power Off Count: 0
Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-01
Module Name: 1.254.A
Mac Address: 00:a0:98:32:d6:dc
IOM Type: IOM6
Shelf Serial Number: 6000368103
IP Address: 192.168.2.221
Protocol Version: 1.2.2. 8
Assigner ID: 1.2.2. 8
State: Active
Last Contact: 218
Power Cycle Count: 0
Power Off Count: 0
Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
Module Name: 1.106.A
Mac Address: 00:a0:98:19:53:ee

IOM Type: IOM6E
Shelf Serial Number: SHJMS000000001A
IP Address: 192.168.3.239
Protocol Version: 2.1.1.21
Assigner ID: 2.1.1.21
State: Initializing
Last Contact: 206
Power Cycle Count: 0
Power Off Count: 0
Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
Module Name: 1.106.B
Mac Address: 00:a0:98:19:55:16
IOM Type: IOM6E
Shelf Serial Number: SHJMS000000001A
IP Address: 192.168.1.23
Protocol Version: 2.1.1.21
Assigner ID: 2.1.1.21
State: Initializing
Last Contact: 209
Power Cycle Count: 0
Power Off Count: 0
Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
Module Name: 1.10.B
Mac Address: 00:a0:98:32:d6:ac
IOM Type: IOM6
Shelf Serial Number: 6000368103
IP Address: 192.168.2.173
Protocol Version: 1.2.2.8
Assigner ID: 1.2.2.8
State: Initializing
Last Contact: 217
Power Cycle Count: 0
Power Off Count: 0
Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
Module Name: 1.10.A
Mac Address: 00:a0:98:32:d6:dc
IOM Type: IOM6
Shelf Serial Number: 6000368103
IP Address: 192.168.2.221
Protocol Version: 1.2.2.8
Assigner ID: 1.2.2.8
State: Initializing
Last Contact: 220
Power Cycle Count: 0
Power Off Count: 0

Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0

8 entries were displayed.

```
[[ID19dea49f3630bbef3d1f1e0e130648af]]
= storage shelf drawer show-phy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of PHYs per drawer

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage shelf drawer show-phy` command displays information for
drawer PHYs in the storage system. If no parameters are specified, the
default command displays the following information about PHYs:

* Shelf Name
* Drawer Number
* PHY Number
* Type
* SAS Address
* State

To display detailed information about a single PHY, use the `-shelf` , `-
drawer` , and `-phy` parameters.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all drawer PHYs, in column style output.

| ``[-instance  ]``::
Displays expanded information for all drawer PHYs in the system. If a
shelf, drawer, and PHY are specified, then this parameter displays the
same detailed information for the PHY you specify as does the -shelf,
-drawer, and -phy parameters.

``[-shelf <text>]`` - Shelf Name::
```

```
Displays the PHYs in the storage shelf that matches the specified shelf
name.

``[-drawer <integer>]`` - Drawer Number::
Displays the PHYs in the drawers that match the specified drawer number.

``[-phy <integer>]`` - PHY Number::
Displays the PHYs that match the specified PHY number.

``[-node {<nodename>|local}]`` - Node Name::
Displays the PHYs that are present for the specified node.

``[-type {unknown|disk|virtual|input}]`` - Type::
Displays the PHYs with the specified type.

``[-physical-id <integer>]`` - Physical ID::
Displays the PHYs that match the specified physical-id.

``[-sas-address <text>]`` - Attached SAS Address::
Displays the PHYs with the specified attached sas address.

``[-state-a {unknown|enabled|disabled}]`` - State Module A::
Displays the PHYs for which module A has the specified state.

``[-state-b {unknown|enabled|disabled}]`` - State Module B::
Displays the PHYs for which module B has the specified state.

``[-status-a <Drawer PHY Status>]`` - Status Module A::
Displays the PHYs with module A currently operating under the specified
status.

``[-status-b <Drawer PHY Status>]`` - Status Module B::
Displays the PHYs with module B currently operating under the specified
status.

== Examples

The following example displays information about all drawer PHYs:
```

cluster1::> storage shelf drawer show-phy
Shelf Drawer PHY # Type SAS Address PHY State A/B
----- ------ ----- -------- ---------- ----------
2.5
1
0 disk 00c5005079183f85 enabled/enabled
1 disk - enabled/enabled
2 disk - enabled/enabled

3 disk 00c50050e1183f85 enabled/enabled
4 disk - enabled/enabled
5 disk - enabled/enabled
6 disk 00c50050dd183f85 enabled/enabled
7 disk - enabled/enabled
8 disk - enabled/enabled
9 disk 00c500502d163f85 enabled/enabled
10 disk - enabled/enabled
11 disk - enabled/enabled
12 input 80090a5045e46f06 enabled/enabled
13 input 80090a5045e46f06 enabled/enabled
14 input 80090a5045e46f06 enabled/enabled
15 input 80090a5045e46f06 enabled/enabled
16 virtual 8a090a503dd01b17 enabled/enabled
2
0 disk 00c500503d0e3d85 enabled/enabled
1 disk - enabled/enabled
2 disk - enabled/enabled
3 disk 00c50050e9173f85 enabled/enabled
4 disk - enabled/enabled
5 disk - enabled/enabled
6 disk 00c50050a9163f85 enabled/enabled
7 disk - enabled/enabled
8 disk - enabled/enabled
9 disk 00c5005021173f85 enabled/enabled
10 disk - enabled/enabled
11 disk - enabled/enabled
12 input 80090a5045e46f06 enabled/enabled
13 input 80090a5045e46f06 enabled/enabled
14 input 80090a5045e46f06 enabled/enabled
15 input 80090a5045e46f06 enabled/enabled
16 virtual 8a090a503d90fd16 enabled/enabled
3
0 disk 00c500503d163f85 enabled/enabled
1 disk - enabled/enabled
2 disk - enabled/enabled
3 disk 00c50050bd163f85 enabled/enabled
4 disk - enabled/enabled
5 disk - enabled/enabled
6 disk 00c50050c1d44085 enabled/enabled
7 disk - enabled/enabled
8 disk - enabled/enabled
9 disk 00c50050f1d54085 enabled/enabled
10 disk - enabled/enabled
11 disk - enabled/enabled
12 input 80090a5045e46f06 enabled/enabled
13 input 80090a5045e46f06 enabled/enabled
14 input 80090a5045e46f06 enabled/enabled
15 input 80090a5045e46f06 enabled/enabled
16 virtual 8a090a503d202a17 enabled/enabled
4
0 disk 00c50050fdd54085 enabled/enabled
1 disk - enabled/enabled
2 disk - enabled/enabled

3 disk 00c50050d9d44085 enabled/enabled
4 disk a0cc0050e5973712 enabled/enabled
5 disk - enabled/enabled
6 disk 00c500506dd34085 enabled/enabled
7 disk - enabled/enabled
8 disk - enabled/enabled
9 disk 00c5005045d64085 enabled/enabled
10 disk - enabled/enabled
11 disk - enabled/enabled
12 input 80090a5045e46f06 enabled/enabled
13 input 80090a5045e46f06 enabled/enabled
14 input 80090a5045e46f06 enabled/enabled
15 input 80090a5045e46f06 enabled/enabled
16 virtual 8a090a503d100b17 enabled/enabled
5
0 disk 00c50050c9d54085 enabled/enabled
1 disk - enabled/enabled
2 disk - enabled/enabled
3 disk 00c50050f9d44085 enabled/enabled
4 disk - enabled/enabled
5 disk - enabled/enabled
6 disk 00c5005081d34085 enabled/enabled
7 disk - enabled/enabled
8 disk - enabled/enabled
9 disk 00c500505dd64085 enabled/enabled
10 disk - enabled/enabled
11 disk - enabled/enabled
12 input 80090a5045e46f06 enabled/enabled
13 input 80090a5045e46f06 enabled/enabled
14 input 80090a5045e46f06 enabled/enabled
15 input 80090a5045e46f06 enabled/enabled
16 virtual 8a090a503df00a17 enabled/enabled
85 entries were displayed.
cluster1::>

```
The following example displays expanded information for PHY 0 of drawer 1
in shelf 2.5:
```

cluster1::> storage shelf drawer show-phy -shelf 2.5 -drawer 1 -phy 0
Shelf: 2.5
Drawer ID: 1
PHY Number: 0
Type: disk
Physical ID: 1
SAS Address: 00c5005079183f85
State A: enabled
State B: enabled
Status A: enabled-12gbs
Status B: enabled-12gbs
cluster1::>

```
[[ID4236b966a99646ba0f6bbb99e7b7cf87]]
= storage shelf drawer show-slot
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a map between bay number and drawer/slot number

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage shelf drawer show-slot` command maps each drawer and slot
number to the corresponding bay number.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields in column style output.

| ``[-instance  ]``::
Displays all slot information.

``[-shelf <text>]`` - Shelf Name::
Displays the slots in the shelf that matches the specified shelf name.

``[-bay <integer>]`` - Bay Number::
Displays the slots that have the specified bay number.

``[-node {<nodename>|local}]`` - Node Name::
Displays the slots that are present for the specified node.

``[-drawer <integer>]`` - Drawer Number::
Displays the slots in the drawers that match the specified drawer number.

``[-slot <integer>]`` - Slot Number::
Displays the slots that match the specified slot number.

``[-is-installed {yes|no}]`` - Is Disk Installed::
Displays the slots that have a disk installed.

== Examples

The following example displays the mapping from drawer and slot number to
```

```
bay number:
```

cluster1::> storage shelf drawer show-slot
Shelf Drawer Slot Bay Installed?
----- ------ ---- --- ----------
2.5
1
0 0 yes
1 1 no
2 2 no
3 3 yes
4 4 no
5 5 no
6 6 yes
7 7 no
8 8 no
9 9 yes
10 10 no
11 11 no
2
0 12 yes
1 13 no
2 14 no
3 15 yes
4 16 no
5 17 no
6 18 yes
7 19 no
8 20 no
9 21 yes
10 22 no
11 23 no
3
0 24 yes
1 25 no
2 26 no
3 27 yes
4 28 no
5 29 no
6 30 yes
7 31 no
8 32 no
9 33 yes
10 34 no
11 35 no
4
0 36 yes
1 37 no
2 38 no
3 39 yes
4 40 yes
5 41 no
6 42 yes

7 43 no
8 44 no
9 45 yes
10 46 no
11 47 no
5
0 48 yes
1 49 no
2 50 no
3 51 yes
4 52 no
5 53 no
6 54 yes
7 55 no
8 56 no
9 57 yes
10 58 no
11 59 no
60 entries were displayed.
cluster1::>

```
[[IDf01454b321452bf96268aea99167176c]]
= storage shelf drawer show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of drawers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage shelf drawer show` command displays information for storage
shelf drawers in the storage system. If no parameters are specified, the
default command displays the following information for the drawers:

* Shelf Name
* Drawer Number
* Status
* Closed/Open
* Disk Count
* Firmware

To display detailed information for a single drawer, use the `-shelf` and
`-drawer` parameters.
```

```
== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all drawers, in column style output.

| ``[-errors  ]``::
Displays the following error status information about the drawers that
have errors:
+
  * Status
  * Error Description



| ``[-instance  ]``::
Displays expanded information for all drawers in the system. If a shelf
and drawer are specified, then this parameter displays the same detailed
information for the specified drawer as does the -shelf and -drawer
parameters.

``[-shelf <text>]`` - Shelf Name::
Displays the drawers in the storage shelf that matches the specified shelf
name.

``[-drawer <integer>]`` - Drawer Number::
Displays the drawers that match the specified drawer number.

``[-node {<nodename>|local}]`` - Node Name::
Displays the drawers that are present for the specified node.

``[-disk-count <integer>]`` - Drawer Disk Count::
Displays the drawers that have the specified disk count.

``[-part-number <text>]`` - Part Number::
Displays the drawers that have the specified part number.

``[-serial-number <text>]`` - Serial Number::
Displays the drawer that matches the specified serial number.

``[-is-closed {open|closed}]`` - Drawer is Closed?::
Displays the drawers that are closed or open.

``[-firmware-a <text>]`` - Firmware A::
Displays the the drawers for which module A has the specified firmware
version.
```

```
``[-firmware-b <text>]`` - Firmware B::
Displays the drawers for which module B has the specified firmware
version.

``[-path-a {unknown|ok|degraded|none}]`` - Path A::
Displays the drawers for which module A has the specified path status.

``[-path-b {unknown|ok|degraded|none}]`` - Path B::
Displays about drawers for which module B has the specified path status.

``[-is-supported {yes|no}]`` - Drawer is Supported?::
Displays the drawers that are supported (TRUE) or not supported (FALSE).

``[-vendor <text>]`` - Vendor Name::
Displays the drawers that match the specified vendor.

``[-mfg-date <text>]`` - Mfg. Date::
Displays the drawers that match the specified manufactured date.

``[-fru-type <text>]`` - FRU Type::
Displays the drawers that match the specified FRU type.

``[-status-a {unknown|normal|warning|error|critical}]`` - Status A::
Displays the drawers with module A currently operating under the specified
status.

``[-status-b {unknown|normal|warning|error|critical}]`` - Status B::
Displays the drawers with module B currently operating under the specified
status.

``[-error <text>]`` - Error::
Displays the drawers that match the specified error description.

== Examples

The following example displays information about all drawers:
```

cluster1::> storage shelf drawer show
Drawer Disk
Shelf Drawer Status A/B Closed? Count Firmware A/B
----- ------ ---------- ------- ----- ----------
2.5
1 normal/normal closed 4 00000634/00000634
2 normal/normal closed 4 00000634/00000634
3 normal/normal closed 4 00000634/00000634
4 normal/normal closed 5 00000634/00000634
5 normal/normal closed 4 00000634/00000634
5 entries were displayed.

cluster1::>

```
The following example displays expanded information about drawer 1 in
shelf 2.5:
```

cluster1::> storage shelf drawer show -shelf 2.5 -drawer 1
Shelf: 2.5
Drawer ID: 1
Part Numer: 111-03071
Serial Number: 021604008153
Drawer is Closed?: closed
Disk Count: 4
Firmware A: 00000634
Firmware B: 00000634
Path A: ok
Path B: ok
Status A: normal
Status B: normal
Drawer is Supported?: yes
Vendor Name: NETAPP
Mfg. Date: 02/2016
FRU Type: SASDRWR
Error Description: -
cluster1::>

```
The following example displays error information about the drawers that
have errors:
```

cluster1::> storage shelf drawer show -errors
Shelf Drawer Status A/B Error Description
----- ------ ---------- ----------
2.5
2 warning/warning Drawer open.
cluster1::>

```
[[IDe3b014fd4a9dfe385c251eebe64ceb96]]
= storage shelf firmware show-update-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Shelf Firmware Update (SFU) Status.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.
```

```
== Description

The `storage shelf firmware show-update-status` command displays the state
of the Shelf Firmware Update process.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - Node [.advaced-priv]#(privilege: advanced)#::
Selects the node that matches this parameter value.

``[-update-status {running|idle}]`` - Disk Shelf Firmware Update Status
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes whose SFU process status matches this parameter value.
Possible values are:
+
  * running - Disk shelf firmware update is in progress.
  * idle - Disk shelf firmware update is not in progress.



``[-in-progress-count <integer>]`` - Number of Shelves with Earlier
Revisions Being Updated [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that matches the number of shelves the SFU process is
updating to this parameter value. This specifies the number of shelves
with earlier revisions that are being updated.

== Examples
```

cluster1::*> storage shelf firmware show-update-status
Update In-Progress
Node Status Count
---------- ------- ----------
cluster-n1 running 10
cluster-n2 idle -
cluster-n3 running 7

```
[[ID9576118625fbaf827b595777c09a42bf]]
= storage shelf firmware update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Update Shelf Firmware

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage shelf firmware update` command updates the firmware on one or
more shelves. You can download the latest firmware by using the
xref:{relative_path}storage-firmware-download.html[storage firmware
download] command. You can specify a shelf whose firmware is to be updated
by using the `-shelf` parameter. You can update the firmware on all the
shelves by not providing the `-shelf` parameter. All the shelves of a
specific module type can be updated by providing a value to the `-module-
type` parameter.

== Parameters

{ ``[-shelf <text>]`` - Shelf Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the shelf whose firmware is to be updated.

| ``[-module-type
{atfcx|esh4|iom3|iom6|iom6e|iom12|iom12e|iom12b|nsm100|nsm8e|iom12g|nsm16e
|iom12c|nsm100b}]`` - Shelf Module Type [.advaced-priv]#(privilege:
advanced)# }::
Update the firmware on the shelves that match the module-type you specify.

``[-refresh <true>]`` - Refresh [.advaced-priv]#(privilege: advanced)#::
Forces an update on the shelf with the highest revision of the applicable
firmware, resulting in a refresh of the firmware image already present on
the shelf.

== Examples

The following example updates the firmware on all the shelves in the
cluster:
```

cluster1::*> storage shelf firmware update

```
The following example updates the firmware on all shelves with the IOM6
module type:
```

cluster1::*> storage shelf firmware update -module-type IOM6

```
The following example updates the firmware on shelf 1.2:
```

cluster1::*> storage shelf firmware update -shelf 1.2

```
The following example refreshes the firmware on all shelves with the IOM6
module type:
```

cluster1::*> storage shelf firmware update -refresh -module-type IOM6

```
The following example refreshes the firmware on shelf 1.2:
```

cluster1::*> storage shelf firmware update -refresh -shelf 1.2

```
== Related Links

* xref:{relative_path}storage-firmware-download.html[storage firmware
download]




[[ID0a557648ec345866962b83d36a4e1e0d]]
= storage shelf location-led modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the state of the shelf Location LED

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage shelf location-led modify` command modifies the on/off state
of the shelf location LED.

== Parameters

| ``-shelf-name <text>`` - Shelf Name }::
This parameter specifies the shelf whose LED is to be turned on or turned
off.

``[-led-status {on|off}]`` - Location LED::
This parameter specifies whether the shelf location LED needs to be turned
on or turned off.

== Examples

The following example turns on the shelf location LED of the specified
shelf.
```

cluster1::> storage shelf location-led modify -shelf-name 1.0 -led-status on

Info: Shelf locate request successful for shelf "1.0".

```
The following example turns off the shelf location LED of the specified
shelf.
```

cluster1::> storage shelf location-led modify -shelf-name 1.0 -led-status off

Info: Shelf locate request successful for shelf "1.0".

```
[[IDda1ada815d65cf3a1a3010d5cfa52767]]
= storage shelf location-led show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Location LED status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage shelf location-led show` command displays the state of shelf
location LED.

== 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.

``[-shelf-name <text>]`` - Shelf Name::
Selects the shelves whose shelf-name matches this parameter value.

``[-node {<nodename>|local}]`` - Node Name::
Selects the nodes that match this parameter value.

``[-stack-id <integer>]`` - Stack ID::
Selects the shelves whose stack-id matches this parameter value.
```

```
``[-shelf-id <integer>]`` - Shelf ID::
Selects the shelves whose shelf-id matches this parameter value.

``[-led-status {on|off}]`` - Location LED::
Shows the state of the shelf location LED.

== Examples

The following example shows the state of the shelf location LED for each
shelf.
```

cluster1::> storage shelf location-led show
Shelf Name Stack ID Shelf ID LED Status
---------- -------- -------- -----------
8.2 8 2 off
8.3 8 3 off
6.0 6 0 unsupported
8.1 8 1 off
4 entries were displayed.

```
[[ID8d15f7e295cfdcb7bcab5d49dfb98f66]]
= storage shelf port show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display storage shelf ports

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage shelf port show` command displays information for storage
shelf ports in the storage system. If no parameters are specified, the
default command displays the following information for the ports:

* Shelf Name
* ID
* Module
* State
* Internal?

To display detailed information for a single port, use the `-shelf` and `-
id` parameters.
```

```
== Parameters

{ ``[-fields <fieldname>,...]``::
Displays output in column style about the specified fields for all shelf
ports.

| ``[-cables  ]``::
Displays information about all cables connected to the shelf ports.

| ``[-instance  ]``::
Displays expanded information for all shelf ports in the system. If a
shelf and ID are specified, then this parameter displays the same detailed
information for the specified port as does the -shelf and -id parameters.

``[-shelf <text>]`` - Shelf Name::
Displays the ports in the storage shelf that matches the specified shelf
name.

``[-id <integer>]`` - Port ID::
Displays the ports that match the specified ID.

``[-node {<nodename>|local}]`` - Node Name::
Displays the ports that are present for the specified node.

``[-module-id {A|B}]`` - Module ID::
Displays the ports from the specified shelf module ID.

``[-is-internal {true|false}]`` - Is Port Internal?::
Displays the ports that are internal.

``[-location <text>]`` - Location::
Displays the ports with the specified location.

``[-is-cable-connected {true|false}]`` - Is Cable Connected?::
Displays the ports that have cables connected to them.

``[-is-error {true|false}]`` - Any Errors?::
Displays the ports for which errors have been logged.

``[-connector-state {connected|disconnected|error}]`` - Connector State::
Displays the ports with the specified connector state.

``[-connector-serial-number <text>]`` - Connector Serial Number::
Displays the ports with the specified connector serial number.

``[-connector-type {QSFP|QSFP+|QSFP28|Mini-SAS-HD}]`` - Connector Type::
```

```
Displays the ports with the specified connector type.

``[-cable-vendor <text>]`` - Cable Vendor::
Displays the ports that are connected to a cable from the specified
vendor.

``[-cable-part-number <text>]`` - Cable Part Number::
Displays the ports that are connected to a cable with the specified part
number.

``[-cable-technology {active-copper|passive-copper|optical}]`` - Cable
Technology::
Displays the ports that are connected to a cable with the specified
technology.

``[-cable-length <text>]`` - Cable Length::
Displays the ports that are connected to a cable with the specified
length.

``[-cable-id <text>]`` - Cable ID::
Displays the ports that are connected to a cable with the specified ID.

``[-cable-end {end_0|end_1}]`` - Cable End::
Displays the ports that are connected to a cable with the specified cable
end.

``[-designator <text>]`` - Designator::
Displays the ports with the specified designator.

``[-wwn <text>]`` - Local Device WWN::
Displays the ports with the specified WorldWide Name (WWN).

``[-remote-wwn <text>]`` - Remote Device WWN::
Displays the ports connected to the specified remote WorldWide Name (WWN).

``[-remote-phy <text>]`` - Remote Phy::
Displays the ports connected to the specified remote PHY.

``[-swap-count <integer>]`` - Swap Count::
Displays the ports with the specified swap count.

``[-mac <MAC Address>]`` - Local MAC Address::
Displays the ports with the specified MAC address.

``[-remote-mac <MAC Address>]`` - Remote MAC Address::
Displays the ports connected to the specified MAC address.
```

```
``[-remote-port <text>]`` - Remote Port::
Displays the ports connected to the specified port.

``[-remote-chassis <text>]`` - Remote Chassis::
Displays the ports connected to the specified chassis.

``[-remote-device <text>]`` - Remote Device::
Displays the ports connected to the specified device.

``[-vlan-id <integer>]`` - VLAN ID::
Displays the ports with the specified Virtual LAN (VLAN) ID.

``[-link-state {unknown|online|offline}]`` - Link State::
Displays the ports with the specified link state.

== Examples

The following example displays information about all shelf ports:
```

cluster1::> storage shelf port show

Shelf ID Module State Internal?
----- — ------ ---------- ---------
1.4
0 A connected false
1 A connected false
2 B connected false
3 B connected false
4 entries were displayed.

```
The following example displays expanded information about port 0 in shelf 1.4:
```

cluster1::> storage shelf port show -shelf 1.4 -id 0
Shelf Name: 1.4
Port ID: 0
Module ID: A
Is Port Internal?: false
Location: rear of the shelf at the top left, on shelf module A
Is Cable Connected?: true
Any Errors?: false
Connector State: connected
Connector Serial Number: 616930439
Connector Type: qsfp+
Cable Vendor: Molex Inc.
Cable Part Number: 112-00431+A0
Cable Technology: passive-copper
Cable Length: 5m
Cable ID: 500a0980000b6c3f-50000d1703544b80

Cable End: end_1
Designator: sqr
Local Device WWN: 500A0980000B6C3F
Remote Device WWN: 50000D1703544B80
Remote Phy: 12
Swap Count: 0

```
The following example displays information about the cables:
```

cluster1::> storage shelf port show -cables

Shelf: 1.4

ID Vendor Part Number Technology Length Type Serial Number — ---------- ---------- ---------- -------
---------- ----------
0 Molex Inc. 112-00431+A0 passive-copper 5m qsfp+ 616930439
1 Molex Inc. 112-00431+A0 passive-copper 5m qsfp+ 616930364
2 Molex Inc. 112-00431+A0 passive-copper 5m qsfp+ 616930452
3 Molex Inc. 112-00431+A0 passive-copper 5m qsfp+ 616930474
4 entries were displayed.

```
:leveloffset: -1


= storage stackmon commands

:leveloffset: +1


[[ID3bab960942eaa0c25adefce0da019fd8]]
= storage stackmon modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]


*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage stackmon modify` command modifies stackmon.

== Parameters
```

```
``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node on which to modify stackmon.

``[-paused {true|false}]`` - Paused [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies whether stackmon is paused. If set to true,
stackmon will not respond to any topology changes.

``[-reset-timer <true>]`` - Reset Stackmon Timer [.advaced-
priv]#(privilege: advanced)#::
This parameter resets stackmon's timer, which will trigger it to check for
and resolve any stack ID conflicts.

``[-refresh-shelf-table <true>]`` - Refresh Shelf Table [.advaced-
priv]#(privilege: advanced)#::
This parameter refreshes the table used by `storage shelf` commands.

== Examples

The following example pauses stackmon on all nodes.
```

cluster1::> storage stackmon modify -node * -paused true

```
The following example triggers stackmon on all nodes.
```

cluster1::> storage stackmon modify -node * -reset-timer true -paused false

```
[[IDaa32dc83bb1ac5fbfb43e6ee0967c02b]]
= storage stackmon show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display stackmon information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage stackmon show` command displays information for stackmon on
all nodes. If no parameters are specified, the default command displays
```

```
the following information:

* Node
* State
* Paused

To display detailed information, use the `-node` parameter.

== Parameters

``[-fields <fieldname>,...]``::
Displays output in column style about the specified fields.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Displays expanded information for stackmon on all nodes. If a node is
specified, then this parameter displays the same detailed information for
the specified port as does the -node parameter.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Displays the stackmon information for the the specified node.

``[-state {init|processing|stable|failed-reconcile}]`` - State [.advaced-
priv]#(privilege: advanced)#::
Displays the stackmon information that matches the specified state.

``[-paused {true|false}]`` - Paused [.advaced-priv]#(privilege:
advanced)#::
Displays the stackmon information that matches the specified value.

``[-timer-delay <integer>]`` - Timer Delay [.advaced-priv]#(privilege:
advanced)#::
Displays the stackmon information that matches the specified timer-delay
value.

== Examples

The following example displays stackmon information for all nodes:
```

cluster1::> storage stackmon show

Node State Paused
--- --- ---
node1 stable false
node2 stable false
node3 stable false
node4 stable false

4 entries were displayed.

```
[[ID03471db17bc276a1e9ba843c99c6023e]]
= storage stackmon repair run
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Repair stackmon

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage stackmon repair run` command diagnoses stackmon
reconciliation failures and attempts to repair the stack IDs.

== Parameters

``[-skip-check-stackmon-state <true>]`` - Skip Check for Stackmon State
[.advaced-priv]#(privilege: advanced)#::
This parameter allows the command to be executed when stackmon is still
running.

== Examples

The following example diagnoses and repairs the stack IDs following a
stackmon reconciliation failure.
```

cluster1::*> storage stackmon repair run

```
[[IDdf2082c2f212ee32ee3b7800c73fa4df]]
= storage stackmon repair show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display repairs that will be made to stackmon if \"storage stackmon repair
run\" is executed
```

```
*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage stackmon repair show` command diagnoses stackmon
reconciliation failures and displays the repairs that will be made if
xref:{relative_path}storage-stackmon-repair-run.html[storage stackmon
repair run] is executed:

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-shelf-uid <integer>]`` - Shelf UID [.advaced-priv]#(privilege:
advanced)#::
Displays the stackmon repairs for the specified shelf. Use
xref:{relative_path}storage-shelf-show.html[storage shelf show -fields
shelf-uuid] to find the shelf-uid.

``[-old-stack-id <integer>]`` - Old Stack ID [.advaced-priv]#(privilege:
advanced)#::
Displays the stackmon repairs with the specified current stack-id.

``[-new-stack-id <integer>]`` - New Stack ID [.advaced-priv]#(privilege:
advanced)#::
Displays the stackmon repairs with the specified new stack-id.

``[-skip-check-stackmon-state <true>]`` - Skip Check for Stackmon State
[.advaced-priv]#(privilege: advanced)#::
This parameter allows the command to be executed when stackmon is still
running.

== Examples

The following example displays the stackmon repairs:
```

cluster 1::*> storage stackmon repair show

Shelf Old Stack New Stack
--- --- ---
123370712932616784 1 1
292812051842665040 3 1
4342319975102154064 2 2
4558492757215937872 2 2
5350844816656434512 1 1
9274403680772754000 1 1
9346461274810681936 2 2
10786768730639108688 2 2
8 entries were displayed.

```
== Related Links

* xref:{relative_path}storage-stackmon-repair-run.html[storage stackmon
repair run]
* xref:{relative_path}storage-shelf-show.html[storage shelf show]




[[ID3377c1cc6eac07aba4a1c64ed31ce41d]]
= storage stackmon topology show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display stackmon topology

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage stackmon topology show` command displays the shelf topology
reported by stackmon, SES, and the disk driver on all nodes. Note: all
node/shelf pairs are shown, even if the node is not connected to the
shelf. The chain-handle, ses-channels, and disk-channels fields will be
empty ("-") when the node is not connected to the shelf.

== Parameters

``[-fields <fieldname>,...]``::
Displays output in column style about the specified fields.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
```

```
Displays expanded information for stackmon topology on all nodes. If a
node and shelf-uid are specified, then this parameter displays the same
detailed information for the specified port as does the -node and -shelf
-uid parameters.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Displays the topology information for the specified node.

``[-shelf-uid <integer>]`` - Shelf UID [.advaced-priv]#(privilege:
advanced)#::
Displays the topology information for the specified shelf. Use
xref:{relative_path}storage-shelf-show.html[storage shelf show]`-
fields```_shelf-uuid_`` to find the shelf-uid.

``[-stack-id <integer>]`` - Stack ID [.advaced-priv]#(privilege:
advanced)#::
Displays the topology information for shelves with the specified stack-id.

``[-chain-handle <integer>]`` - Stackmon Chain Handle [.advaced-
priv]#(privilege: advanced)#::
Displays the topology information for shelves with the specified chain-
handle.

``[-reconcile-state {processing|succeeded|failed}]`` - Reconcile State
[.advaced-priv]#(privilege: advanced)#::
Displays the topology information for shelves with the specified
reconcile-state.

``[-ses-channels <text>,...]`` - SES Channels [.advaced-priv]#(privilege:
advanced)#::
Displays the topology information for shelves with the specified SES
channels.

``[-disk-channels <text>,...]`` - Disk Channels [.advaced-
priv]#(privilege: advanced)#::
Displays the topology information for shelves with the specified disk
channels.

``[-ses-disk-match {true|false}]`` - SES Matches Disk Topology [.advaced-
priv]#(privilege: advanced)#::
Displays the topology information for shelves that are seen on the same
channels by SES and the disk driver.

== Examples

The following example displays topology information for all shelves:
```

```
cluster1::*> storage stackmon topology show
Failed
Node Shelf Chain Handle Stack Reconcile
---- -------------------- -------------------- ----- ---------
node1
123370712932616784 1152922178327261751 1 false
292812051842665040 1152922178327261751 1 false
4342319975102154064 1152922178327258740 2 false
4558492757215937872 1152922178327258740 2 false
5350844816656434512 1152922178327261751 1 false
9274403680772754000 1152922178327261751 1 false
9346461274810681936 1152922178327258740 2 false
10786768730639108688 1152922178327258740 2 false
node2
123370712932616784 1152922178315916511 1 false
292812051842665040 1152922178315916511 1 false
4342319975102154064 1152922178315913225 2 false
4558492757215937872 1152922178315913225 2 false
5350844816656434512 1152922178315916511 1 false
9274403680772754000 1152922178315916511 1 false
9346461274810681936 1152922178315913225 2 false
10786768730639108688 1152922178315913225 2 false
node3
123370712932616784 1152922178329970138 1 false
292812051842665040 1152922178329970138 1 false
4342319975102154064 1152922178329977389 2 false
4558492757215937872 1152922178329977389 2 false
5350844816656434512 1152922178329970138 1 false
9274403680772754000 1152922178329970138 1 false
9346461274810681936 1152922178329977389 2 false
10786768730639108688 1152922178329977389 2 false
node4
123370712932616784 1152922178327647260 1 false
292812051842665040 1152922178327647260 1 false
4342319975102154064 1152922178327645097 2 false
4558492757215937872 1152922178327645097 2 false
5350844816656434512 1152922178327647260 1 false
9274403680772754000 1152922178327647260 1 false
9346461274810681936 1152922178327645097 2 false
10786768730639108688 1152922178327645097 2 false
32 entries were displayed.
```

```
== Related Links

* xref:{relative_path}storage-shelf-show.html[storage shelf show]




:leveloffset: -1
```

```
= storage tape commands

:leveloffset: +1


[[ID061cd35e3684f39ca5a5fc51e25d9e0a]]
= storage tape offline
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Take a tape drive offline

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command takes the specified tape drive offline.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node to which the tape drive is
attached.

{ ``-name <text>`` - Tape Drive Device Name::
Use this parameter to specify the device name of the tape drive that needs
to be taken offline. The format of the device `-name` name includes a
prefix to specify how the tape cartridge is handled and a suffix to
describe the density of the tape. The prefix suggests 'r', 'nr' or 'ur'
for rewind, no rewind, or unload/reload and a suffix shows density of 'l',
'm', 'h' or 'a'. For example, a tape device name for this operation might
have the form "nrst8m" were 'nr' is the 'no rewind' prefix, 'st8' is the
alias-name and 'm' is the tape density. You can use the 'storage tape show
-device-names' command to find more information about device names of tape
drives attached to a node.

| ``-device-id <text>`` - Tape Drive Device ID }::
 Use this parameter to specify the device ID of the tape drive that needs
to be taken offline.

== Examples

The following example takes the tape drive with device name 'nrst8m'
```

```
offline. This tape drive is attached to cluster1-01.
```

cluster1::> storage tape offline -node cluster1-01 -name nrst8m

```
[[ID4e53feb1bff6990738f45af80e4caa86]]
= storage tape online
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Bring a tape drive online

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command brings a specified tape drive online.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node to which the tape drive is
attached.

{ ``-device-id <text>`` - Tape Drive Device ID::
 Use this parameter to specify the device ID of the tape drive that needs
to be brought online.

| ``-name <text>`` - Tape Drive Device Name }::
Use this parameter to specify the device name of the tape drive that needs
to be brought online. The format of the device `-name` name includes a
prefix to specify how the tape cartridge is handled and a suffix to
describe the density of the tape. The prefix suggests 'r', 'nr' or 'ur'
for rewind, no rewind, or unload/reload and a suffix shows density of 'l',
'm', 'h' or 'a'. For example, a tape device name for this operation might
have the form "nrst8m" were 'nr' is the 'no rewind' prefix, 'st8' is the
alias-name and 'm' is the tape density. You can use the 'storage tape show
-device-names' command to find more information about device names of tape
drives attached to a node.

== Examples

The following example brings the tape drive with device id sw4:2.126L4
attached to the node, cluster1-01, online.
```

cluster1::> storage tape online -node cluster1-01 -device-id sw4:2.126L4

```
[[ID0ad1c2a1492482906ef11a0e87d9e8cb]]
= storage tape position
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a tape drive cartridge position

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command changes the tape drive cartridge position.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node to which the tape drive is
attached.

``-name <text>`` - Tape Drive Device Name::
Use this parameter to specify the device name of the tape drive whose
cartridge position needs to be changed. The format of the device `-name`
includes a prefix to specify how the tape cartridge is handled and a
suffix to describe the density of the tape. The prefix suggests 'r', 'nr'
or 'ur' for rewind, no rewind, or unload/reload and a suffix shows density
of 'l', 'm', 'h' or 'a'. For example, a tape device name for this
operation might have the form "nrst8m" were 'nr' is the 'no rewind'
prefix, 'st8' is the alias-name and 'm' is the tape density. You can use
the 'storage tape show -device-names' command to find more information
about device names of tape drives attached to a node.

``-operation {weof|fsf|bsf|fsr|bsr|rewind|erase|eom}`` - Tape Position
Operation::
Use this parameter to specify the tape positioning operation. The possible
values for `-operation` are:
+
  * weof - Write end-of-file marks
  * fsf - Forward space end-of-file marks
  * bsf - Backward space end-of-file marks
  * fsr - Forward space records
  * bsr - Backward space records
  * rewind - Rewind the tape
```

```
  * erase - Erase then entire tape media from current position
  * eom - Position the tape at end of data (end of media if full)



``[-count <integer>]`` - Count for Positioning::
Use this parameter to specify the count for a tape positioning operation.
You can specify this parameter only with the following operations: weof,
fsf, bsf, fsr, and bsr. The default value of this parameter is one.

== Examples

The following example specifies a rewind operation on a tape device. Note
the -count parameter does not need to be specified for this type of
operation.
```

cluster1::> storage tape position -node cluster1-01 -name nrst8m -operation rewind

```
The following example specifies an fsf (forward space filemark) operation
on a tape device. Note the -count parameter specifies 5 forward space
filemarks for this operation.
```

cluster1::> storage tape position -node cluster1-01 -name nrst1a -operation fsf -count 5

```
The following example specifies an eom (end-of-media) operation on a tape
device. The 'eom' positions a tape at end of data (end of media if full).
Note the -count parameter does not need to be specified for this type of
operation.
```

cluster1::> storage tape position -node cluster1-01 -name rst0h -operation eom

```
[[ID3e3fafa8b97f3ace3e226ad9d5fab7f9]]
= storage tape reset
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset a tape drive

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command resets a specified tape drive.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node to which the tape drive is
attached.

``-device-id <text>`` - Tape Drive Device ID [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the device ID of the tape drive is to be
reset.

== Examples

The following example resets the tape drive with device ID sw4:2.126L3
attached to the node, cluster1-01.
```

cluster1::> storage tape reset -node cluster1-01 -device-id sw4:2.126L3

```
[[IDd34548f2bb6efc1f92721272efc6a9a1]]
= storage tape show-errors
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display tape drive errors
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage tape show-errors` command displays error information about
tape drives attached to the cluster. By default, this command displays the
following information about all tape drives:

* Node to which the tape drive is attached
* Device ID of the tape drive
* Type of device(tape drive)
* Description of the tape drive
* Alias name of the tape drive
* Tape drive errors



== 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>|local}]`` - Node::
Displays detailed information about tape drives on the specified node.

``[-device-id <text>]`` - Device ID::
Selects the tape drive with the specified device ID.

``[-device-type <text>]`` - Device Type::
Selects the devices with the specified type of tape drive.

``[-description <text>]`` - Description::
Selects the tape drives with the specified description.

``[-alias-name <text>]`` - Alias Name::
Selects the tape drive with the specified alias name.

``[-wwnn <text>]`` - World Wide Node Name::
Selects the tape drives with the specified World Wide Node Name.
```

```
``[-wwpn <text>]`` - World Wide Port Name::
Selects the tape drive with the specified World Wide Port Name.

``[-serial-number <text>]`` - Serial Number::
Selects the tape drive with the specified serial number.

``[-error <text>]`` - Tape Drive Error Description::
Selects the tape drives with the specified error string.

``[-initiator <text>]`` - Initiator Port::
Selects the tape drives with the specified initiator port.

== Examples

The following example shows error information for all tape drives attached
to cluster1.
```

cluster1::> storage tape show-errors
Node: node1
Device ID: 0d.125
Device Type: tape drive
Description: Hewlett-Packard LTO-5
Alias: st0
Errors: hardware error; repair or replace tape drive
Node: node1
Device ID: 2d.0
Device Type: tape drive
Description: IBM LTO-6 ULT3580
Alias: st2
Errors: -

```
The following example shows error information for tape drive sw4:2.126L1
attached to the node, node1.
```

cluster1::> storage tape show-errors -device-id sw4:2.126L1 -node node1
Node: node1
Device ID: sw4:2.126L1
Device Type: tape drive
Description: Hewlett-Packard LTO-3
Alias: st3
Errors: -

```
[[ID943dd29bd5dcc3f1a1a4bb37be33c3f7]]
= storage tape show-media-changer
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about media changers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This `storage tape show-media-changer` command displays information about
media changers attached to the cluster. By default, this command displays
the following information about all media changers:

* Device ID of media changer
* Description of media changer
* World Wide Node Name of media changer
* World Wide Port Name of media changer
* Serial number of media changer
* Media changer errors
* Node to which the media changer is attached
* Initiator port which hosts the media changer
* Alias name of media changer
* Operational state of media changer
* Functional status of media changer



== 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.

``[-device-id <text>]`` - Device ID::
Selects the media changer with the specified device ID.

``[-node {<nodename>|local}]`` - Node::
Displays detailed information about media changers on the specified node.

``[-description <text>]`` - Description::
```

```
Selects the media changers with the specified description.

``[-alias-name <text>]`` - Alias Name::
Selects the media changer with the specified alias name.

``[-wwnn <text>]`` - World Wide Node Name::
Selects the media changers with the specified World Wide Node Name.

``[-wwpn <text>]`` - World Wide Port Name::
>Selects the media changer with the specified World Wide Port Name.

``[-serial-number <text>]`` - Serial Number::
Selects the media changer with the specified serial number.

``[-device-if-type {unknown|fibre-channel|SAS|pSCSI}]`` - Device If Type::
Selects the media changers with the specified interface type.

``[-device-state {unknown|available|ready-write-enabled|ready-write-
protected|offline|in-use|error|reserved-by-another-host|normal}]`` -
Operational State of Device::
Selects the media changers with the specified operational state.

``[-error <text>]`` - Media Changer Error Description::
Selects the media changers with the specified error string.

``[-initiator <text>]`` - Initiator Port::
Selects the media changers with the specified initiator port.

== Examples

The following example displays information about all media changers
attached to the cluster:
```

cluster1::> storage tape show-media-changer

Media Changer: sw4:10.11L1
Description: PX70-TL
WWNN: 2:00a:000e11:10b919
WWPN: 2:00b:000e11:10b919
Serial Number: 00FRU7800000_LL1
Errors: -

Paths:
Node Initiator Alias Device State Status
---------- --------- ------- ---------- --------
cluster1-01 2b mc0 in-use normal

Media Changer: sw4:12.4L1

Description: NEO-TL
WWNN: 2:001:000e11:10b919
WWPN: 2:002:000e11:10b919
Serial Number: 00FRU7800000_LL0
Errors: -
Paths:
Node Initiator Alias Device State Status
---------- --------- ------- ---------- --------
cluster1-01 5a mc1 available normal

```
[[ID604bf02dd3e8b5a09c4986c8ad962f74]]
= storage tape show-supported-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays the qualification and supported status of tape drives

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the supported and qualification status of all tape
drives recognized by ONTAP attached to a node in the cluster. This
includes nonqualified tape drives. Such tape drives do not have a Tape
Configuration File (TCF) on the storage system. A nonqualified tape drive
can be used if the tape drive emulates a qualified tape drive or if the
appropriate TCF for the nonqualified tape drive is downloaded from the
NetApp Support Site to the storage system.

== 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>|local}]`` - Node::
Selects the tape drives that match this parameter value.
```

```
``[-tape-drive <text>]`` - Tape Drive Name::
Selects the tape drives that match this parameter value.

``[-is-supported {true|false}]`` - Tape Drive Supported::
Selects the tape drives that match this parameter value.

``[-status <text>]`` - Supported Status::
Selects the tape drives that match this parameter value.

== Examples

The following example displays support and qualification status of tape
drives recognized by ONTAP. The command also identifies tape drives
attached to the node that are nonqualified (not supported).
```

cluster1::> storage tape show-supported-status
Node: Node1
Is
Tape Drive Supported Support Status
---------- --------- ----------
sw4:2.126L6 false Nonqualified tape drive
Hewlett-Packard C1533A true Qualified
Hewlett-Packard C1553A true Qualified
Hewlett-Packard Ultrium 1 true Qualified
Sony SDX-300C true Qualified
Sony SDX-500C true Qualified
StorageTek T9840C true Dynamically Qualified
StorageTek T9840D true Dynamically Qualified
Tandberg LTO-2 HH true Dynamically Qualified
The following example displays support and qualification status of tape
drives selected by tape-drive. The command identifies the supported
status of the selected tape drive.
cluster1::> storage tape show-supported-status -tape-drive "Sony SDX-300C"
Node: Node1
Is
Tape Drives Supported Support Status
---------- --------- ----------
Sony SDX-300C true Qualified

```
[[ID5d0ecceba58b523813666eb920d437b5]]
= storage tape show-tape-drive
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about tape drives
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This `storage tape show-tape-drive` command displays information about
tape drives attached to the cluster. By default, this command displays the
following information about all tape drives:

* Device ID of tape drive
* Description of tape drive
* World Wide Node Name of tape drive
* World Wide Port Name of tape drive
* Serial Number of tape drive
* Tape drive errors
* Node to which the tape drive is attached
* Initiator port which hosts the tape drive
* Alias name of tape drive
* Operational state of tape drive
* Functional status of tape drive



== 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.

``[-device-id <text>]`` - Device ID::
Selects the tape drive with the specified device ID.

``[-node {<nodename>|local}]`` - Node::
Displays detailed information about tape drives on the specified node.

``[-description <text>]`` - Description::
Selects the tape drives with the specified description.

``[-alias-name <text>]`` - Alias Name::
Selects the tape drive with the specified alias name.

``[-wwnn <text>]`` - World Wide Node Name::
```

```
Selects the tape drives with the specified World Wide Node Name.

``[-wwpn <text>]`` - World Wide Port Name::
Selects the tape drive with the specified World Wide Port Name.

``[-serial-number <text>]`` - Serial Number::
Selects the tape drive with the specified serial number.

``[-device-if-type {unknown|fibre-channel|SAS|pSCSI}]`` - Device If Type::
Selects the tape drives with the specified interface type.

``[-device-state {unknown|available|ready-write-enabled|ready-write-
protected|offline|in-use|error|reserved-by-another-host|normal}]`` -
Operational State of Device::
Selects the tape drives with the specified operational state.

``[-error <text>]`` - Tape Drive Error Description::
Selects the tape drives with the specified error string.

``[-initiator <text>]`` - Initiator Port::
Selects the tape drives with the specified initiator port.

``[-resv-type {off|persistent|scsi}]`` - Reservation type for device::
Selects the tape drives with the specified type.

== Examples

The following example displays information about all tape drives attached
to the cluster:
```

```
cluster1::> storage tape show-tape-drive
Tape Drive: sw4:11.126
Description: StorageTek T10000C
WWNN: 5:001:04f000:b39ec8
WWPN: 5:001:04f000:b39ec9
Serial Number: 576004000041
Errors: -

Paths:
Node Initiator Alias Device State Status
----------------------- --------- ------- ----------------------- --------
cluster1-01 2a st0 ready-write-enabled normal
Tape Drive: sw4:12.4
Description: HP LTO-3
WWNN: 2:001:000e11:10b919
WWPN: 2:002:000e11:10b919
Serial Number: 1068000371
Errors: -
```

Paths:
Node Initiator Alias Device State Status
---------- --------- ------- ---------- --------
cluster1-01 0b st1 ready-write-enabled normal

```
[[ID1d6f082c3e4be12a4da4713c2c56bb4d]]
= storage tape show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about tape drives and media changers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage tape show` command displays information about tape drives and
media changers attached to the cluster. Where it appears in the remainder
of this document "device" may refer to either a tape drive or a media
changer. By default, this command displays the following information about
all tape drives and media changers:

* Node to which the tape drive/media changer is attached
* Device ID of the tape drive/media changer
* Description of the tape drive/media changer
* Type of device: tape drive or media changer
* Functional status of the tape drive/media changer



== 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.

| ``[-alias  ]``::
Displays the tape drive/media changer alias with the following details:
+
  * Node to which tape drive/media changer is attached
  * Device ID of the tape drive/media changer
  * Alias name of the tape drive/media changer
```

```
  * Alias mapping for tape drive/media changer



| ``[-connectivity  ]``::
Displays the connectivity from the node to the tape drive/media changer
with the following details:
+
  * Node to which tape drive/media changer is attached
  * Device ID of the tape drive/media changer
  * Tape drive/media changer description
  * Type of device: tape drive or media changer
  * Interface type for the tape drive/media changer
  * World Wide Node Name of tape drive/media changer
  * World Wide Port Name of tape drive/media changer
  * Serial Number of tape drive/media changer
  * Tape drive/media changer errors
  * Initiator port which hosts the tape drive/media changer
  * Alias name of the tape drive/media changer
  * Operational state of tape drive/media changer
  * Functional status of tape drive/media changer



| ``[-device-names  ]``::
Displays the tape drive names for used tape positioning using the
following details: rewind, no rewind, unload/reload and density
+
  * Node to which tape drive/media changer is attached
  * Device ID of the tape drive/media changer
  * Tape drive/media changer description
  * Device Names that include Rewind, no Rewind, Unload/Reload



| ``[-status  ]``::
Displays the status of tape drive/media changer with the following
details:
+
  * Device ID of the tape drive/media changer
  * Tape drive/media changer description
  * World Wide Node Name of tape drive/media changer
  * World Wide Port Name of tape drive/media changer
  * Serial Number of tape drive/media changer
  * Alias name of the tape drive/media changer
  * Format used for tape cartridge mounted by tape drive
```

```
  * Tape drive/media changer errors
  * Node to which tape drive/media changer is attached
  * Operational state of tape drive/media changer
  * File number following last tape drive I/O operation
  * Block number following last tape drive I/O operation
  * Residual count following last tape drive I/O operation



| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-device-id <text>]`` - Device ID::
Selects the tape drive/media changer with the specified device ID.

``[-node {<nodename>|local}]`` - Node::
Displays detailed information about tape drives or media changers on the
specified node.

``[-device-type <text>]`` - Device Type::
Selects the devices with the specified type of tape drive or media
changer.

``[-description <text>]`` - Description::
Selects the tape drives/media changers with the specified description.

``[-alias-name <text>]`` - Alias Name::
Selects the tape drive/media changer with the specified alias name.

``[-alias-mapping <text>]`` - Alias Mapping::
Selects the tape drive/media changer with the specified alias mapping.

``[-wwnn <text>]`` - World Wide Node Name::
Selects the tape drives/media changers with the specified World Wide Node
Name.

``[-wwpn <text>]`` - World Wide Port Name::
Selects the tape drive/media changer with the specified World Wide Port
Name.

``[-serial-number <text>]`` - Serial Number::
Selects the tape drive/media changer with the specified serial number.

``[-functional-status {unknown|normal|error}]`` - Functional Status::
Selects the tape drives/media changers with the specified functional
status of the device.
```

```
``[-device-if-type {unknown|fibre-channel|SAS|pSCSI}]`` - Device Interface
Type::
Selects the tape drives/media changers with the specified interface type.

``[-device-state {unknown|available|ready-write-enabled|ready-write-
protected|offline|in-use|error|reserved-by-another-host|normal}]`` -
Operational State of Device::
Selects the tape drives/media changers with the specified operational
state.

``[-format <text>,...]`` - Tape Cartridge Format::
Selects the tape drives with the specified tape format.

``[-error <text>]`` - Tape Error::
Selects the tape drives/media changers with the specified error string.

``[-initiator <text>]`` - Initiator Port::
Selects the tape drives/media changers with the specified initiator port.

``[-file-number <integer>]`` - File Number::
Selects the tape drives/media changers with the specified file number. The
file number is the number of file marks between the beginning of media and
current logical position. File number gets modified on write file mark,
and forward or backward space file operations. A value of -1 indicates
unknown position on the tape media or tape not loaded in the tape drive.

``[-block-number <integer>]`` - Block Number::
Selects the tape drives/media changers with the specified block number.
The block number is the number of logical blocks between the beginning of
tape media or the prior file mark and the current logical position on the
tape media. Block number gets modified on writes, reads, and forward or
backward space over records (blocks). The block number also gets reset to
zero when a file mark is crossed or another file mark is written that
designates a new file. If the tape is back spaced to a prior file mark,
the block number might be zeroed. A value of -1 indicates unknown position
on the tape media or that a tape not loaded in the tape drive.

``[-residual-count <integer>]`` - Residual Count of Last I/O Operation::
Selects the tape drives with the specified residual count.

``[-device-name-r <text>,...]`` - Device Name for Rewind::
Selects the tape drives with the specified device name for rewind.

``[-device-name-nr <text>,...]`` - Device Name for No Rewind::
Selects the tape drives with the specified device name for no rewind.
```

```
``[-device-name-ur <text>,...]`` - Device Name for Unload Reload::
Selects the tape drives with the specified device name for unload/reload.

``[-resv-type {off|persistent|scsi}]`` - Reservation Type for device::
Selects the tape drives with the specified type.

``[-aliases-name <text>,...]`` - Alias Names::
Selects the tape drive/media changer with the specified alias name.

``[-aliases-mapping <text>,...]`` - Alias Mappings::
Selects the tape drive/media changer with the specified alias mapping.

== Examples

The following example displays information about all tape drives and media
changers attached to the cluster:
```

cluster1::> storage tape show

Node: cluster1-01
Device ID Device Type Description Status
--- --- --- ---
sw4:10.11 tape drive HP LTO-3 error
Node: cluster1-01
Device ID Device Type Description Status
--- --- --- ---
sw4:10.11L1 media changer PX70-TL normal

```
The following example displays detailed information about a tape drive
named sw4:10.11
```

cluster1::> storage tape show -device-id sw4:10.11

Node: cluster1-01
Device ID Device Type Description Status
--- --- --- ---
sw4:10.11 tape drive HP LTO-3 error

```
[[ID7e637cfcad9e4dad7e493637b5f34493]]
= storage tape trace
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable/disable tape trace operations

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command enables or disables diagnostic tape trace operations for all
tape drives attached to the node you have specified.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node on which the tape trace feature is
enabled or disabled.

``[-is-trace-enabled {true|false}]`` - Tape Trace Enabled or Disabled::
Use this parameter to enable or disable the tape trace feature. By
default, the tape trace feature is enabled.

== Examples

The following example enables tape trace operation on the node, cluster1-
01.
```

cluster1::> storage tape trace -node cluster1-01 -is-trace-enabled true

```
[[ID505c0c74712072ad49f0342807246d18]]
= storage tape alias clear
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear alias names
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command clears alias names for a tape drive or media changer.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node to which the tape drive is
attached.

{ ``-name <text>`` - Alias Name That Is to Be Cleared::
Use this parameter to specify the alias name that is to be cleared. You
can use the 'storage tape show -alias' command to find more information
about alias names of tape drives and media changers attached to a node.
The `-clear-scope` and `-name` parameters are mutually exclusive. If you
specify the `-name` parameter, a single alias name is cleared.

| ``-clear-scope {tape|media-changer|all}`` - Scope of Alias Clear
Operation }::
Use this parameter to specify the scope of the alias clear operation. The
`-clear-scope` and `-name` parameters are mutually exclusive. If you
specify the `-clear-scope` parameter, multiple aliases are cleared
depending upon the value of the parameter.
+
The possible values for `-clear-scope` are as follows:
+
  * tape - Clear all tape drive aliases
  * media-changer - Clear all media-changer aliases
  * all - Clear both tape drive and media-changer aliases



== Examples

The following example clears an alias name 'st3' attached to the node,
cluster1-01.
```

cluster1::> storage tape alias clear -node cluster1-01 -name st3

```
The following example clears all tape drive alias names attached to the
node, cluster1-01.
```

cluster1::> storage tape alias clear -node cluster1-01 -clear-scope tape

```
The following example clears all media changer alias names attached to the
node, cluster1-01.
```

cluster1::> storage tape alias clear -node cluster1-01 -clear-scope media-changer

```
The following example clears both tape and media changer alias names
attached to the node, cluster1-01.
```

cluster1::> storage tape alias clear -node cluster1-01 -clear-scope all

```
[[IDdf8ca378f5f082a277cfacd5b9c1cef9]]
= storage tape alias set
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set an alias name for tape drive or media changer

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command sets an alias name for a tape drive or media changer.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node to which the tape drive is
attached.

``-name <text>`` - Alias Name for Tape Drive or Media Changer::
Use this parameter to specify the alias name for tape drive or media
changer. For a tape drive alias name, the format is 'st' followed by one
or more digits. For a media changer alias name, the format is 'mc'
followed by one or more digits.

``-mapping <text>`` - Mapping for Alias Name::
Use this parameter to specify the mapping for an alias name. Use the
format 'SN[``_<serial-number>_`` ]'. Valid mapping for serial numbers are
in the format 'SN[``_<serial-number>_`` ]' where the ``_<serial-number>_``
is from 2 to 124 characters long and includes the following characters: 0-
9, a-z, and A-Z.

== Examples

The following example sets an alias name 'st3' for a tape drive with
serial number SN[123456]L4 attached to the node, node1.
```

cluster1::storage tape alias> set -node node1 -name st3 -mapping SN[123456]L4.

```
The following example sets an alias name 'mc1' for a media changer with
serial number SN[65432] attached to the node, node1.
```

cluster1::storage tape alias> set -node node1 -name mc1 -mapping SN[65432].

```
[[ID9a6670180960072a89f865f7cd4937e4]]
= storage tape alias show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays aliases of all tape drives and media changers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays aliases of all tape drives and media changers
attached to every node in the cluster.

== 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>|local}]`` - Node::
Displays the tape drive and media changer aliases on the specified node.

``[-alias-list <text>,...]`` - List of Aliases::
Displays the node with the specified list of tape drive and media changer
aliases.

== Examples

The following example shows the aliases of all tape drives and media
changers attached to every node in the cluster:
```

cluster1::> storage tape alias show

Node: node1

Alias Mapping

```
------------------------------------ ----------------------------------------
mc0 SN[00FRU7800000_LL0]L1
mc1 SN[00FRU7800000_LL1]L1
mc2 SN[aa6a64c69360a0980248c8]
mc3 SN[c940abe8b0c3a0980248c8]
mc4 SN[fba082e6b335a0980248c8]L5
st0 SN[HU19487T7N]
st1 SN[1068000230]
st10 SN[fba0c508b335a0980248c8]L7

Node: node2

Alias Mapping
------------------------------------ ----------------------------------------
mc1 SN[c940982fc48c8]
st3 SN[ST456HT8N]L3
st2 SN[HG68000230]L2
st11 SN[aba673980248c8]L7
```

```
[[ID05fb8c724be6889c54cf6ad1f3759f69]]
= storage tape config-file delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a tape config file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `storage tape config-file delete` command deletes the specified tape
drive configuration file from all nodes that are currently part of the
cluster.

== Parameters

``-filename <text>`` - Config File Filename [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the tape configuration file that will
be deleted from all nodes that are currently part of the cluster.

== Examples

The following example deletes the specified tape drive configuration files
on every node that is currently part of the cluster:
```

cluster1::> storage tape config-file delete -filename XYZ_LTO-6.TCF

```
[[ID703de4dd726f2569eca02b33672984b3]]
= storage tape config-file get
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Get a tape drive configuration file

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage tape config-file get` command uploads a specified tape drive
configuration file to each node that is currently part of the cluster.

== Parameters

``-url <text>`` - Config File URL::
This parameter specifies the URL that provides the location of the package
to be fetched. Standard URL schemes, including HTTP and TFTP, are
accepted.

== Examples

The following example uploads the specified tape drive configuration file
to each node that is currently part of the cluster:
```

cluster1::> storage tape config-file get -url http://example.com/~tapeconfigfile/XYZ_LTO-6.TCF

```
[[IDe115b5fb47c7ba2b3958f9e37e0c4365]]
= storage tape config-file show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of tape drive configuration files on the given node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage tape config-file show` command lists the tape drive
configuration files loaded onto each node in the cluster.

== 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>|local}]`` - Node::
Selects information about tape drive configuration files for the specified
node.

``[-config-file <text>]`` - Tape Config File::
Selects information about the tape drive configuration file specified.

== Examples

The following example lists the tape drive config files loaded onto each
node in the cluster:
```

cluster1::> storage tape config-file show

Node: node1

Tape Config Files

CERTANCE_LTO2_ULTRIUM.TCF
CERTANCE_LTO3_ULTRIUM.TCF
HP_LT09.TCF
HP_LTO2.TCF
HP_LTO3_ULTRIUM.TCF
HP_LTO4_ULTRIUM.TCF
HP_LTO5_ULTRIUM.TCF
HP_LTO6_ULTRIUM.TCF
IBM_3592.TCF
IBM_3592E05.TCF
IBM_5038_sdfkjl.TCF
IBM_LTO2_ULT3580.TCF
IBM_LTO2_ULTRIUM.TCF

```
[[IDd69b324917936572e46de8880cdec30f]]
= storage tape library config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display connectivity to back-end storage tape libraries.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays information such as how the storage tape libraries
connect to the cluster, LUN groups, number of LUNs, WWPN, and switch port
information. Use this command to verify the cluster's storage tape library
configuration or to assist in troubleshooting.

== 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.

| ``[-switch  ]``::
If you specify this parameter, switch port information is shown.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-node {<nodename>|local}]`` - Controller Name::
The name of the clustered node for which information is being displayed.

``[-group <integer>]`` - LUN Group::
A LUN group is a set of LUNs that shares the same path set.

``[-target-wwpn <text>]`` - Library Target Ports::
The World Wide Port Name of a storage tape library port.

``[-initiator <text>]`` - Initiator::
The host bus adapter that the clustered node uses to connect to storage
tape libraries.

``[-array-name <array name>]`` - Library Name::
Name of the storage tape library that is connected to the clustered node.

``[-target-side-switch-port <text>]`` - Target Side Switch Port::
This identifies the switch port that connects to the tape library's target
port.

``[-initiator-side-switch-port <text>]`` - Initiator Side Switch Port::
This identifies the switch port that connects to the node's initiator
port.

``[-lun-count <integer>]`` - Number of  LUNS::
This is a command-line switch (-lun-count) used to restrict what LUN
groups are displayed in the output.

== Examples

The following example displays the storage tape library configuration
information.
```

cluster1::> storage tape library config show
LUN LUN
Node Group Count Library Name Library Target Port Initiator
--- --- --- --- --- ---
cluster1-01
0 2 TAPE_LIB_1 50050763124b4d6f 3d

cluster1::>

```
[[IDf4122469cb71bfa9e8755e8e3e7f0965]]
= storage tape library path show-by-initiator
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of LUNs on the given tape library

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays path information for every initiator port connected
to a tape library. The output is similar to the storage library path show
command but the output is listed by initiator.

== Parameters

{ ``[-fields <fieldname>,...]``::
fields used to be used in this display

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Controller name::
The name of the clustered node for which information is being displayed.

``[-initiator <text>]`` - Initiator Port::
Initiator port that the clustered node uses.

``[-target-wwpn <text>]`` - Target Port::
Target World Wide Port Name. Port on the storage tape library that is
being used.

``[-initiator-side-switch-port <text>]`` - Initiator Side Switch Port::
Switch port connected to the clustered node.

``[-target-side-switch-port <text>]`` - Target Side Switch Port::
Switch port connected to the tape library.

``[-array-name <array name>]`` - Library Name::
Name of the storage tape library that is connected to the cluster.

``[-tpgn <integer>]`` - Target Port Group Number::
TPGN refers to the target port group to which the target port belongs. A
target port group is a set of target ports which share the same LUN access
characteristics and failover behaviors.
```

```
``[-port-speed <text>]`` - Port Speed::
Port Speed of the specified port.

``[-path-io-kbps <integer>]`` - Kbytes of I/O per second on Path (Rolling
Average)::
Rolling average of Kbytes of I/O per second on the library path.

``[-path-iops <integer>]`` - Number of I/O per second on Path (Rolling
Average)::
Rolling average of I/O per second on the library path.

``[-initiator-io-kbps <integer>]`` - Kbytes of I/O per second on Initiator
(Rolling Average)::
Rolling average of Kbytes of I/O per second on the initiator port.

``[-initiator-iops <integer>]`` - Number of I/O per second on Initiator
(Rolling Average)::
>Rolling average of I/O per second on the initiator port.

``[-target-io-kbps <integer>]`` - Kbytes of I/O per second to Target
(Rolling Average)::
Rolling average of Kbytes of I/O per second on the target port.

``[-target-iops <integer>]`` - Number of I/O per second to Target (Rolling
Average)::
Rolling average of I/O per second on the target port.

== Examples

The following example displays the path information by initiator for a
storage tape library.
```

cluster1::> storage tape library path show-by-initiator
Node: cluster1-01
Initiator I/O Initiator Side Path I/O Target Side Target I/O
Initiator (KB/s) Switch Port (KB/s) Switch Port (KB/s) Target Port Library Name
--- --- --- --- --- --- --- ---
0b 0 sw_tape:6 0 sw_tape:0 0 510a09800000412d TAPE_LIB_1
sw_tape:1 0 510a09820000412d TAPE_LIB_1
3d 0 N/A 0 N/A 0 50050763124b4d6f TAPE_LIB_2
3 entries were displayed.

```
[[ID7a44bef4d21b7127f2919a9f87fe7577]]
= storage tape library path show
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of Tape Libraries on the given path

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays path information for a tape library and has the
following parameters by default:

* Node name
* Initiator port
* Target port
* TPGN (Target Port Group Number)
* Port speeds
* Path I/O in Kbytes/sec
* IOPs



== Parameters

{ ``[-fields <fieldname>,...]``::
fields used to be used in this display

| ``[-detail  ]``::
Using this option displays the following:
+
  * Target IOPs
  * Target LUNs
  * Path IOPs
  * Path errors
  * Path quality
  * Path LUNs
  * Initiator IOPs
  * Initiator LUNs



| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-node {<nodename>|local}]`` - Controller name::
The name of the clustered node for which information is being displayed.

``[-array-name <array name>]`` - Library Name::
Name of the storage tape library that is connected to the cluster.

``[-target-wwpn <text>]`` - Target Port::
Target World Wide Port Name. Port on the storage tape library that is
being used.

``[-initiator <text>]`` - Initiator Port::
Initiator port that the clustered node uses.

``[-initiator-side-switch-port <text>]`` - Initiator Side Switch Port::
Switch port connected to the clustered node.

``[-tpgn <integer>]`` - Target Port Group Number::
TPGN refers to the target port group to which the target port belongs. A
target port group is a set of target ports which share the same LUN access
characteristics and failover behaviors.

``[-port-speed <text>]`` - Port Speed::
Port Speed of the specified port.

``[-path-io-kbps <integer>]`` - Kbytes of I/O per second on Path(Rolling
Average)::
Rolling average of Kbytes of I/O per second on the library path.

``[-path-iops <integer>]`` - Number of I/O per second on Path(Rolling
Average)::
Rolling average of I/O per second on the library path.

``[-initiator-io-kbps <integer>]`` - Kbytes of I/O per second on
Initiator(Rolling Average)::
Rolling average of Kbytes of I/O per second on the initiator port.

``[-initiator-iops <integer>]`` - Number of I/O per second on
Initiator(Rolling Average)::
>Rolling average of I/O per second on the initiator port.

``[-target-io-kbps <integer>]`` - Kbytes of I/O per second to
Target(Rolling Average)::
Rolling average of Kbytes of I/O per second on the target port.

``[-target-iops <integer>]`` - Number of I/O per second to Target(Rolling
Average)::
Rolling average of I/O per second on the target port.
```

```
``[-target-side-switch-port <text>]`` - Target Side Switch Port::
Switch port connected to the tape library.

``[-path-link-errors <integer>]`` - Link Error count on path::
Fibre Channel link error count.

``[-path-quality <integer>]`` - Percentage of weighted error threshold::
A number representing the threshold of errors that is allowed on the path.
Path quality is a weighted error value. When the error weight of a path
exceeds the threshold, I/O is routed to a different path.

``[-path-lun-in-use-count <integer>]`` - Number of LUNs in the in-use
state on this path::
Number of LUNs on this path.

``[-initiator-lun-in-use-count <integer>]`` - Number of LUNs in the in-use
state on this initiator::
Number of LUNs on this initiator.

``[-target-lun-in-use-count <integer>]`` - Number of LUNs in the in-use
state on this target::
Number of LUNs on this target.

== Examples

The following example displays the path information for a storage tape
library
```

```
cluster1::> storage tape library path show
Node Initiator Target Port TPGN Speed (KB/s) IOPs
-------------------- --------- ----------------------- ------ ------- ------------ ------------
cluster1-01
3d 50050763124b4d6f 61 4 Gb/S 0 0
cluster1-01
0b 510a09800000412d 35 4 Gb/S 0 0
cluster1-01
0b 510a09820000412d 1 4 Gb/S 0 0
3 entries were displayed.
```

```
[[ID7f38ba08888c563ba70173b706ef4528]]
= storage tape load-balance modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the tape load balance configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage tape load-balance modify` command modifies the tape load
balance setting for a specified node in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node on which the tape load balance setting
is to be modified.

``[-is-enabled {true|false}]`` - Is Tape Load Balance Enabled::
This parameter specifies whether tape load balancing is enabled on the
node. The default setting is false.

== Examples

The following example modifies the tape load balance setting on node1 in
the cluster:
```

cluster1::> storage tape load-balance modify -node node1 -is-enabled true

```
[[IDdb098e1030415a56e851a32c312dc51f]]
= storage tape load-balance show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays the tape load balance configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage tape load-balance show` command displays tape load balance
settings for each node in the cluster.

== 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>|local}]`` - Node::
Selects information about tape load balancing for the specified node.

``[-is-enabled {true|false}]`` - Is Tape Load Balance Enabled::
Selects information about load balance configuration as specified by
enabled or disabled setting.

== Examples

The following example shows the load balance setting for each node in the
cluster:
```

cluster1::> storage tape load-balance show

Node Enabled
--- ---
node1 false

node2 false
2 entries were displayed.

```
:leveloffset: -1


:leveloffset: -1


= system commands

:leveloffset: +1


= system bridge commands

:leveloffset: +1


[[ID4eb21a7788fd215e187bbf84a2d390e6]]
= system bridge add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a bridge for monitoring

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system bridge add` command enables you to add FC-to-SAS bridges for
in-band monitoring in a MetroCluster configuration.

== Parameters

``-name <text>`` - Bridge Name::
This parameter identifies the bridge being added for in-band monitoring.

== Examples

The following command adds a bridge with name ATTO_FibreBridge7600N_1 for
monitoring:
```

cluster1::> system bridge add -name ATTO_FibreBridge7600N_1
cluster1::> stor br show
Symbolic Is Monitor Managed
Bridge Name Monitored Status By Vendor Model Bridge WWN
---------- --------- --------- ------- ------- ------ ---------- ----------
ATTO_2000001086a183c0 bridge1
true ok in-band Atto FibreBridge 7600N 2000001086a183c0
ATTO_FibreBridge7600N_2 -
false - none Atto - 2000001086a18340
ATTO_FibreBridge7600N_3 -
false - none Atto - 2000001086a1ee40
ATTO_FibreBridge7600N_4 -
false - none Atto - 2000001086700a30
4 entries were displayed.
cluster1::>

```
[[ID92f0665cf31d677d33b385e6c8270021]]
= system bridge modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a bridge's configuration information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system bridge modify` enables you to modify certain parameters for
identifying and accessing the FC-to-SAS bridges added for monitoring in a
MetroCluster configuration.

== Parameters

``-name <text>`` - Bridge Name::
This parameter specifies the name of the bridge.

``[-node-visible-list <text>,...]`` - Nodes Bridge is Visble To::
This parameter specifies nodes that are visible to the bridge.

== Examples

The following command modifies 'ATTO_2000001086b0' bridge node-visible-
list to 'nodeA,nodeB'
```

cluster1::> system bridge modify -name ATTO_2000001086b0 -node-visible-list nodeA,nodeB
cluster1::>

```
[[IDf59a5d75e02ddf844ef63cadd9b53c43]]
= system bridge refresh
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Refresh bridge info

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system bridge refresh` command triggers a refresh of the SNMP data
for the MetroCluster FC switches and FC-to-SAS bridges. There is no action
if the refresh is already in progress. For systems running ONTAP 9.7 or
earlier, the FC switches and FC-to-SAS bridges must have been previously
added for monitoring by using the the `system switch fiber-channel add`
and xref:{relative_path}system-bridge-add.html[system bridge add] commands
respectively. For systems running ONTAP 9.8 or later, the `system bridge
refresh` command performs the additional task of automatically adding any
bridges that have recently been added to the environment (assuming the
automatic option is set).

== Examples

The following command triggers a refresh for the SNMP data:
```

cluster1::*> system bridge refresh
cluster1::*>

```
== Related Links

* xref:{relative_path}system-bridge-add.html[system bridge add]




[[ID3f85b05af509225b8a89168d26814d03]]
= system bridge remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a bridge from monitoring

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system bridge remove` enables you to remove FC-to-SAS bridges that
were previously added for SNMP monitoring.

== Parameters

``-name <text>`` - Bridge Name::
This parameter specifies the name of the bridge added for monitoring.

== Examples

The following command removes 'ATTO_10.226.197.16' bridge from monitoring:
```

cluster1::> system bridge remove -name ATTO_10.226.197.16
cluster1::> system bridge show
Is Monitor
Bridge Symbolic Name Vendor Model Bridge WWN Monitored Status
---------- ------------- ------- --------- ---------------- --------- -------
ATTO_FibreBridge6500N_1
Bridge Number 16
Atto FibreBridge 6500N
2000001086603824 false -
ATTO_FibreBridge6500N_2
Not Set Atto FibreBridge 6500N
20000010866037e8 false -
ATTO_FibreBridge6500N_3

Not Set Atto FibreBridge 6500N
2000001086609e0e false -
ATTO_FibreBridge6500N_4
Not Set Atto FibreBridge 6500N
2000001086609c06 false -
4 entries were displayed.

```
[[ID02b717218084bf42d59dbb4d506c00f4]]
= system bridge run-cli
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Execute a CLI command on a bridge

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `storage bridge run-cli` command enables you to execute an ATTO bridge
command.

== Parameters

``-name <text>`` - Bridge Name::
This parameter specifies the name of the bridge that the command is to be
executed on.

``-command <text>`` - CLI command to execute::
This parameter specifies the command to be executed on the named bridge.

== Examples

The following example executes a command on a bridge
```

sti8040mcc-201_siteA::> storage bridge run-cli -name ATTO_FibreBridge7500N_1 -command "Help"

```
[[ID418c2a47b29d5f756b5bd9d6de238a30]]
= system bridge show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display bridge information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system bridge show` command displays information about all the
storage bridges in the MetroCluster configuration. The bridges must have
been previously added for monitoring using the xref:{relative_path}system-
bridge-add.html[system bridge add] command. If no parameters are
specified, the default command displays the following information about
the bridges:

* Bridge
* Symbolic Name
* Vendor
* Model
* Bridge WWN
* Is Monitored
* Is Bridge Secure
* Managed By
* Monitor Status

To display detailed profile information about a single bridge, use the `-
name` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all the bridges, in column style output.

| ``[-connectivity  ]``::
Displays the following details about the connectivity from different
entities to the bridge:
+
  * Node
  * Initiator
  * Initiator Side Switch Port
  * Target Side Switch Port
  * Target Port WWN
  * Target Port Number
```

```
| ``[-cooling  ]``::
Displays the following details about the chassis temperature sensor(s) on
the bridge:
+
  * Sensor Name
  * Reading in degree Celsius (C)
  * Fan operational status
  * Minimum Safe Operating Temperature in degree Celsius (C)
  * Maximum Safe Operating Temperature in degree Celsius (C)
  * Sensor Status



| ``[-error  ]``::
Displays the errors related to the bridge.

| ``[-ports  ]``::
Displays the following details about the bridge FC ports:
+
  * Port number
  * Port administrative status
  * Port operational status
  * Port operating mode
  * Port negotiated speed
  * Peer world wide name

+
Displays the following details about the bridge SAS ports:
+
  * Port number
  * Port negotiated data rate
  * Port data rate capability
  * Port PHY1 operational status
  * Port PHY2 operational status
  * Port PHY3 operational status
  * Port PHY4 operational status
  * Port administrative status
  * Port operational status
  * Peer world wide name



| ``[-power  ]``::
Displays the status of the replaceable power supplies for the FibreBridge
7500N only:
+
```

```
  * Power supply name
  * Power supply status



| ``[-sfp  ]``::
Displays the following details about the bridge FC ports Small Form-factor
Pluggable (SFP):
+
  * Port number
  * SFP vendor
  * SFP serial number
  * SFP part number
  * SFP speed capability

+
Displays the following details about the bridge SAS ports Quad Small Form-
factor Pluggable (QSFP):
+
  * Port number
  * QSFP vendor
  * QSFP serial number
  * QSFP type
  * QSFP part number

+
Displays the following details about the bridge SAS ports Mini-SAS HD:
+
  * Port number
  * Mini-SAS HD vendor
  * Mini-SAS HD serial number
  * Mini-SAS HD type
  * Mini-SAS HD part number



| ``[-stats  ]``::
Displays the following details about the bridge FC ports:
+
  * Port number
  * Port operational status
  * Port operational mode
  * Port negotiated speed
  * Port link failure count
  * Port synchronization loss count
  * Port CRC error count
```

```
  * Port operational mode
  * Port received word count (Rx)
  * Port transmitted word count (Tx)

+
Displays the following details about the bridge SAS ports:
+
  * Port number
  * PHY port number
  * Port negotiated speed
  * Port speed capability
  * Port invalid DWORD count
  * Port disparity error count
  * Port synchronization loss count
  * Port PHY reset count
  * Port link changed count
  * Port CRC error count



| ``[-instance  ]`` }::
Displays expanded information about all the bridges in the system. If a
bridge is specified, then this parameter displays the same detailed
information for the bridge you specify as does the -name parameter.

``[-name <text>]`` - Bridge Name::
Displays information only about the bridges that match the name you
specify.

``[-wwn <text>]`` - Bridge World Wide Name::
Displays information only about the bridges that match the bridge WWN you
specify.

``[-model <text>]`` - Bridge Model::
Displays information only about the bridges that match the bridge model
you specify.

``[-vendor {unknown|Atto}]`` - Bridge Vendor::
Displays information only about the bridges that match the bridge vendor
you specify.

``[-fw-version <text>]`` - Bridge Firmware Version::
Displays information only about the bridges that match the bridge firmware
version you specify.

``[-serial-number <text>]`` - Bridge Serial Number::
Displays information only about the bridges that match the bridge serial
```

```
number you specify.

``[-address <IP Address>]`` - Bridge IP Address::
Displays information only about the bridges that match the bridge IP
address you specify.

``[-is-monitoring-enabled {true|false}]`` - Is Monitoring Enabled for
Bridge?::
Displays information only about the bridges that match the bridge
monitoring value you specify.

``[-status {unknown|ok|error}]`` - Bridge Status::
Displays information only about the bridges that match the bridge
monitoring status you specify.

``[-profile-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS
[{+|-}hh:mm]}]`` - Bridge Profile Data Last Successful Refresh Timestamp::
Displays information only about the bridges that match the profile data
last successful refresh timestamp you specify.

``[-symbolic-name <text>]`` - Bridge Symbolic Name::
Displays information only about the bridges that match the symbolic name
you specify.

``[-snmp-community <text>]`` - SNMP Community Set on the Bridge::
Displays information only about the bridges that match the bridge SNMP
community you specify.

``[-managed-by {SNMP|in-band|none}]`` - Bridge Management Method::
This parameter specifies whether the bridge uses the SNMP or in-band
management method. FibreBridge 6500N uses SNMP only; FibreBridge 7500N may
use either.

``[-is-bridge-secure {true|false}]`` - Is Security Enabled For Bridge?::
Displays information only about the bridges that match the bridge security
value you specify.

``[-node-visible-list <text>,...]`` - Nodes Bridge is Visble To::
Displays information only about the bridges that are available to the
node.

``[-error-text-list <text>,...]`` - Bridge Error Description List::
Displays information only about the bridges that have the errors you
specify.

``[-temp-sensor-name <text>]`` - Temperature Sensor Name::
Displays information only about the bridges that have the temperature
```

```
sensor with the name you specify.

``[-min-safe-oper-temp <integer>]`` - Minimum Safe Operating Temperature
in Degree Celsius::
Displays information only about the bridges that have the temperature
sensor with the minimum safe operating temperature you specify.

``[-max-safe-oper-temp <integer>]`` - Maximum Safe Operating Temperature
in Degree Celsius::
Displays information only about the bridges that have the temperature
sensor with the maximum safe operating temperature you specify.

``[-temp-reading <integer>]`` - Chassis Temperature Sensor Reading in
Degree Celsius::
Displays information only about the bridges that have the temperature
sensors with the reading you specify.

``[-temp-sensor-status {normal|warning|critical}]`` - Chassis Temperature
Sensor Status::
Displays information only about the bridges that have the temperature
sensor with the status you specify.

``[-temp-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Bridge Chassis Temperature Data Last Successful Refresh
Timestamp::
Displays information only about the bridges that match the temperature
sensor data last successful refresh timestamp you specify.

``[-fc-port-index-list <integer>,...]`` - Bridge FC Port Index List::
Displays information only about the bridges that have the ports with the
indexes you specify.

``[-fc-port-oper-state-list {unknown|online|offline}]`` - Bridge FC Port
Operational State List::
Displays information only about the bridges that have the ports with the
operational states you specify.

``[-fc-port-admin-state-list {unknown|disabled|enabled}]`` - Bridge FC
Port Admin State List::
Displays information only about the bridges that have the ports with the
administrative states you specify.

``[-fc-port-negotiated-data-rate-list {unknown|2|4|8|16}]`` - Bridge FC
Port Negotiated Data Rate List::
Displays information only about the bridges that have the ports with the
negotiated data rates you specify.
```

```
``[-fc-port-negotiated-conn-mode-list {unknown|loop|n-port}]`` - Bridge FC
Port Negotiated Connection Mode List::
Displays information only about the bridges that have the ports with the
negotiated connection modes you specify.

``[-fc-port-wwn-list <text>,...]`` - Bridge FC Port WWN List::
Displays information only about the bridges that have the ports with the
world-wide names you specify.

``[-fc-port-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS
[{+|-}hh:mm]}]`` - Bridge FC Port Data Last Successful Refresh Timestamp::
Displays information only about the bridges that match the FC ports data
last successful refresh timestamp you specify.

``[-fc-port-stats-index-list <integer>,...]`` - Bridge FC Port Index
List::
Displays information only about the bridges that have the ports with the
indexes you specify.

``[-fc-port-tx-words-list <integer>,...]`` - Bridge FC Port Transmitted
Word Count List::
Displays information only about the bridges that have the ports with the
number of transmitted words you specify.

``[-fc-port-rx-words-list <integer>,...]`` - Bridge FC Port Received Word
Count List::
Displays information only about the bridges that have the ports with the
number of received words you specify.

``[-fc-port-link-failures-list <integer>,...]`` - Bridge FC Port Link
Failure Count List::
Displays information only about the bridges that have the ports with the
number of link failures you specify.

``[-fc-port-sync-losses-list <integer>,...]`` - Bridge FC Port Sync Loss
Count List::
Displays information only about the bridges that have the ports with the
number of synchronization losses you specify.

``[-fc-port-invalid-crc-list <integer>,...]`` - Bridge FC Port Invalid CRC
Count List::
Displays information only about the bridges that have the ports with the
number of invalid CRCs you specify.

``[-fc-port-stats-data-last-successful-refresh-timestamp {MM/DD/YYYY
HH:MM:SS [{+|-}hh:mm]}]`` - Bridge FC Port Stats Last Successful Refresh
Timestamp::
```

```
Displays information only about the bridges that match the FC port stats
data last successful refresh timestamp you specify.

``[-sas-port-index-list <integer>,...]`` - Bridge SAS Port Index List::
Displays information only about the bridges that have the SAS ports with
the indexes you specify.

``[-sas-port-oper-state-list {unknown|online|offline|degraded}]`` - Bridge
SAS Port Operational State List::
Displays information only about the bridges that have the SAS ports with
the operational states you specify.

``[-sas-port-phy1-oper-state-list {unknown|online|offline}]`` - Bridge SAS
Port PHY1 Operational State List::
Displays information only about the bridges that have the SAS ports with
the PHY1 operational states you specify.

``[-sas-port-phy2-oper-state-list {unknown|online|offline}]`` - Bridge SAS
Port PHY2 Operational State List::
Displays information only about the bridges that have the SAS ports with
the PHY2 operational states you specify.

``[-sas-port-phy3-oper-state-list {unknown|online|offline}]`` - Bridge SAS
Port PHY3 Operational State List::
Displays information only about the bridges that have the SAS ports with
the PHY3 operational states you specify.

``[-sas-port-phy4-oper-state-list {unknown|online|offline}]`` - Bridge SAS
Port PHY4 Operational State List::
Displays information only about the bridges that have the SAS ports with
the PHY4 operational states you specify.

``[-sas-port-admin-state-list {unknown|disabled|enabled}]`` - Bridge SAS
Port Administrative State List::
Displays information only about the bridges that have the SAS ports with
the administrative states you specify.

``[-sas-port-data-rate-capability-list
{unknown|1.5Gbps|3Gbps|6Gbps|12Gbps}]`` - Bridge SAS Port Data Rate
Capability List::
Displays information only about the bridges that have the SAS ports with
the data rate capabilities you specify.

``[-sas-port-negotiated-data-rate-list
{unknown|1.5Gbps|3Gbps|6Gbps|12Gbps}]`` - Bridge SAS Port Negotiated Data
Rate List::
Displays information only about the bridges that have the SAS ports with
```

```
the negotiated data rates you specify.

``[-sas-port-wwn-list <text>,...]`` - Bridge SAS Port WWN List::
Displays information only about the bridges that have the SAS ports with
the world-wide names you specify.

``[-sas-port-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS
[{+|-}hh:mm]}]`` - Bridge SAS Port DB Data Last Successful Refresh
Timestamp::
Displays information only about the bridges that match the SAS ports data
last successful refresh timestamp you specify.

``[-sas-port-stats-phy-index-list <integer>,...]`` - Bridge SAS Port PHY
Index List::
Displays information only about the bridges that have the SAS ports with
the PHY indexes you specify.

``[-sas-port-link-changed-list <integer>,...]`` - Bridge SAS Port Link
Changed Count List::
Displays information only about the bridges that have the SAS ports with
the link changed count you specify.

``[-sas-port-invalid-crc-list <integer>,...]`` - Bridge SAS Port Invalid
CRC Count List::
Displays information only about the bridges that have the SAS ports with
the invalid CRCs you specify.

``[-sas-port-phy-reset-list <integer>,...]`` - Bridge SAS Port PHY Reset
Count List::
Displays information only about the bridges that have the SAS ports with
the PHY reset count you specify.

``[-sas-port-sync-losses-list <integer>,...]`` - Bridge SAS Port Sync Loss
Count List::
Displays information only about the bridges that have the SAS ports with
the synchronization losses you specify.

``[-sas-port-disparity-count-list <integer>,...]`` - Bridge SAS Port
Disparity Count List::
Displays information only about the bridges that have the SAS ports with
the disparity count you specify.

``[-sas-port-invalid-dword-list <integer>,...]`` - Bridge SAS Port Invalid
DWORD Count List::
Displays information only about the bridges that have the SAS ports with
the invalid DWORD count you specify.
```

```
``[-sas-port-stats-index-list <integer>,...]`` - Bridge SAS Port Index
List::
Displays information only about the bridges that have the SAS ports with
the indexes you specify.

``[-sas-port-stats-data-rate-capability-list
{unknown|1.5Gbps|3Gbps|6Gbps|12Gbps}]`` - Bridge SAS Port Data Rate
Capability List::
Displays information only about the bridges that have the SAS ports with
the data rate capabilities you specify.

``[-sas-port-stats-negotiated-data-rate-list
{unknown|1.5Gbps|3Gbps|6Gbps|12Gbps}]`` - Bridge SAS Port Negotiated Data
Rate List::
Displays information only about the bridges that have the SAS ports with
the negotiated data rates you specify.

``[-sas-port-stats-data-last-successful-refresh-timestamp {MM/DD/YYYY
HH:MM:SS [{+|-}hh:mm]}]`` - Bridge SAS Port Statistics Data Last
Successful Refresh Timestamp::
Displays information only about the bridges that match the SAS port stats
data last successful refresh timestamp you specify.

``[-fc-sfp-port-index-list <integer>,...]`` - Bridge FC Port Index List::
Displays information only about the bridges that have the FC ports with
the indexes you specify.

``[-fc-port-sfp-vendor-list <text>,...]`` - Bridge FC Port SFP Vendor
List::
Displays information only about the bridges that have the FC ports with
the SFP vendors you specify.

``[-fc-port-sfp-serial-number-list <text>,...]`` - Bridge FC Port SFP
Serial Number List::
Displays information only about the bridges that have the FC ports with
the SFP serial numbers you specify.

``[-fc-port-sfp-part-number-list <text>,...]`` - Bridge FC Port SFP Part
Number List::
Displays information only about the bridges that have the FC ports with
the SFP part numbers you specify.

``[-fc-port-sfp-data-rate-capability-list {2Gb|4Gb|8Gb|16Gb|32Gb}]`` -
Bridge FC Port SFP Data Rate Capability List::
Displays information only about the bridges that have the FC ports with
the SFP data rate capabilities you specify.
```

```
``[-fc-port-sfp-data-last-successful-refresh-timestamp {MM/DD/YYYY
HH:MM:SS [{+|-}hh:mm]}]`` - Bridge FC Port SFP Data Last Successful
Refresh Timestamp::
Displays information only about the bridges that match the FC ports SFP
data last successful refresh timestamp you specify.

``[-sas-qsfp-port-index-list <integer>,...]`` - Bridge SAS Port Index
List::
Displays information only about the bridges that have the SAS ports with
the indexes you specify.

``[-sas-port-qsfp-vendor-list <text>,...]`` - Bridge SAS Port QSFP Vendor
List::
Displays information only about the bridges that have the SAS ports with
the QSFP vendors you specify.

``[-sas-port-qsfp-serial-number-list <text>,...]`` - Bridge SAS Port QSFP
Serial Number List::
Displays information only about the bridges that have the SAS ports with
the QSFP serial numbers you specify.

``[-sas-port-qsfp-type-list {unknown|optical|active-copper|passive-
copper}]`` - Bridge SAS Port QSFP Type List::
Displays information only about the bridges that have the SAS ports with
the QSFP types you specify.

``[-sas-port-qsfp-part-number-list <text>,...]`` - Bridge SAS Port QSFP
Part Number List::
Displays information only about the bridges that have the SAS ports with
the QSFP part numbers you specify.

``[-sas-port-qsfp-data-last-successful-refresh-timestamp {MM/DD/YYYY
HH:MM:SS [{+|-}hh:mm]}]`` - Bridge SAS Port QSFP Data Last Successful
Refresh Timestamp::
Displays information only about the bridges that match the SAS ports QSFP
data last successful refresh timestamp you specify.

``[-mini-sas-hd-index-list <integer>,...]`` - Bridge Mini-SAS HD Index
List::
Displays information only about the bridges that have SAS ports with the
Mini-SAS HD indexes that you specify.

``[-mini-sas-hd-vendor-list <text>,...]`` - Bridge Mini-SAS HD Vendor
List::
Displays information only about the bridges that have SAS ports with the
Mini-SAS HD vendors that you specify.
```

```
``[-mini-sas-hd-serial-number-list <text>,...]`` - Bridge Mini-SAS HD
Serial Number List::
Displays information only about the bridges that have SAS ports with the
Mini-SAS HD serial numbers that you specify.

``[-mini-sas-hd-type-list <text>,...]`` - Bridge Mini-SAS HD Type List::
Displays information only about the bridges that have SAS ports with the
Mini-SAS HD types that you specify.

``[-mini-sas-hd-part-number-list <text>,...]`` - Bridge Mini-SAS HD Part
Number List::
Displays information only about the bridges that have SAS ports with the
Mini-SAS HD part numbers that you specify.

``[-mini-sas-hd-data-last-successful-refresh-timestamp {MM/DD/YYYY
HH:MM:SS [{+|-}hh:mm]}]`` - Bridge Mini-SAS HD Data Last Successful
Refresh Timestamp::
Displays information only about the bridges that match the SAS ports Mini-
SAS HD data with the last successful refresh timestamp that you specify.

``[-power-supply-index-list <integer>,...]`` - Bridge Power Supply Index
List::
Displays information only about the bridges that have power supplies with
the indexes that you specify.

``[-power-supply-name-list <text>,...]`` - Bridge Power Supply Name List::
Displays information only about the bridges that have power supplies with
the name that you specify.

``[-power-supply-status-list {unknown|down|up}]`` - Bridge Power Supply
Status List::
Displays information only about the bridges that have power supplies with
the status that you specify.

``[-power-supply-data-last-successful-refresh-timestamp {MM/DD/YYYY
HH:MM:SS [{+|-}hh:mm]}]`` - Bridge Power Supply Data Last Successful
Refresh Timestamp::
Displays information only about the bridges that match the power supply
last data with the last successful refresh timestamp that you specify.

``[-node-list {<nodename>|local}]`` - Node Name List::
Displays information only about the bridges that are connected to the
nodes you specify.

``[-initiator-list <text>,...]`` - Initiator List::
Displays information only about the bridges that are connected to the
nodes hosting the initiators you specify.
```

```
``[-initiator-side-switch-port-name-list <text>,...]`` - Initiator Side
Switch Port Name List::
Displays information only about the bridges that are connected to the
initiator-side switch ports you specify.

``[-target-side-switch-port-name-list <text>,...]`` - Target Side Switch
Port Name List::
Displays information only about the bridges that are connected to the
target-side switch ports you specify.

``[-target-port-wwn-list <text>,...]`` - Target Port WWN List::
Displays information only about the bridges that match the target ports
with world-wide names you specify.

``[-target-port-index-list <integer>,...]`` - Target Port Index List::
Displays information only about the bridges that match the target ports
with indexes you specify.

== Examples

The following example displays information about all bridges:
```

```
cluster1::> system bridge show
Is Monitor
Bridge Symbolic Name Vendor Model Bridge WWN Monitored Status
---------- ------------- ------- --------- ---------------- --------- -------
ATTO_10.226.197.16
Bridge Number 16 retyped
Atto FibreBridge 6500N
2000001086603824 true ok
ATTO_10.226.197.17
Not Set Atto FibreBridge 6500N
20000010866037e8 true ok
ATTO_10.226.197.18
Not Set Atto FibreBridge 6500N
2000001086609e0e true ok
ATTO_10.226.197.19
Not Set Atto FibreBridge 6500N
2000001086609c06 true ok
4 entries were displayed.
cluster1::>
```

```
The following example displays connectivity (node to bridge) information
about all bridges:
```

```
cluster1::> system bridge show -connectivity
Bridge Name: ATTO_10.226.197.16
```

Bridge WWN: 2000001086603824
Vendor: Atto
Model: FibreBridge 6500N
Serial Number: FB6500N101405
Firmware Version: 1.60 A68E 51.01
Management IP: 10.226.197.16
Errors: -
Initiator Side Target Side Port
Node Initiator Switch Port Switch Port Target Port WWN No
---------- --------- ---------- ---------- ---------- ----
dpg-mcc-3240-15-b1 0c mcc-cisco-8Gb-fab-3:1-29
mcc-cisco-8Gb-fab-1:1-25
2100001086603824 1
dpg-mcc-3240-15-b2 0c mcc-cisco-8Gb-fab-3:1-30
mcc-cisco-8Gb-fab-1:1-25
2100001086603824 1

```
The following command displays cooling (temperature sensors) information
about all bridges:
```

cluster1::> system bridge show -cooling
Bridge Name: ATTO_10.226.197.16
Bridge WWN: 2000001086603824
Vendor: Atto
Model: FibreBridge 6500N
Serial Number: FB6500N101405
Firmware Version: 1.60 A68E 51.01
Management IP: 10.226.197.16
Errors: -
Chassis Temperature Sensor:
Min Safe Max Safe
Sensor Name Reading Oper Temp Oper Temp Status
---------- ------- --------- --------- -------
Chassis 42 0 70 normal
Temperature
Sensor

```
The following command displays the error information about all bridges:
```

cluster1::> system bridge show -error
Bridge Name: ATTO_10.226.197.16
Bridge WWN: 2000001086603824

ATTO_10.226.197.16(2000001086603824):Bridge is Unreachable over Management Network.
Bridge Name: ATTO_10.226.197.17
Bridge WWN: 20000010866037e8

ATTO_10.226.197.17(20000010866037e8):Bridge is Unreachable over Management Network.
Bridge Name: ATTO_10.226.197.18
Bridge WWN: 2000001086609e0e

ATTO_10.226.197.18(2000001086609e0e):Bridge is Unreachable over Management Network.
Bridge Name: ATTO_10.226.197.19
Bridge WWN: 2000001086609c06

ATTO_10.226.197.19(2000001086609c06):Bridge is Unreachable over Management Network.
4 entries were displayed.

```
The following command displays the detailed information about all the
bridges:
```

cluster1::> system bridge show -instance
Bridge Name: ATTO_10.226.197.16
Bridge WWN: 2000001086603824
Vendor: Atto
Model: FibreBridge 6500N
Serial Number: FB6500N101405
Firmware Version: 1.60 A68E 51.01
Management IP: 10.226.197.16
Errors: -

```
The following command displays power supply information about all bridges:
```

cluster1::> system bridge show -power
Bridge Name: ATTO_10.226.197.47
Bridge WWN: 2000001086601506
Vendor: Atto
Model: FibreBridge 6500N
Serial Number: FB6500N100526
Firmware Version: 1.60 069G 51.01
Management IP: 10.226.197.47
Errors: -
Last Update Time: -

```
Bridge Power Supplies:
```

```
Power Supply Name Status
----------------- -------
                - -
```

```
     Bridge Name: ATTO_10.226.197.48
      Bridge WWN: 20000010867002d0
          Vendor: Atto
           Model: FibreBridge 7500N
   Serial Number: FB7500N100018
Firmware Version: 2.00 006U 105.01
   Management IP: 10.226.197.48
          Errors: -
Last Update Time: 10/22/2015 13:37:37 -04:00
```

```
Bridge Power Supplies:
```

```
Power Supply Name Status
----------------- -------
                A up
                B down
```

```
The following command displays port information about all bridges:
```

cluster1::> system bridge show -ports
Bridge Name: ATTO_10.226.197.16
Bridge WWN: 2000001086603824
Vendor: Atto
Model: FibreBridge 6500N
Serial Number: FB6500N101405
Firmware Version: 1.60 A68E 51.01
Management IP: 10.226.197.16
Errors: -
FC Ports:
Admin Oper Neg
Ports Status Status Port Mode Speed WWPN
----- -------- ------- ------------- ------- ----------------
1 enabled online n-port 8gb 2100001086603824
2 enabled offline unknown unknown 2200001086603824
Last Update Time: 8/12/2014 12:34:36 -04:00
SAS Ports:
Neg Data
Data Rate PHY1 PHY2 PHY3 PHY4 Admin Oper
Ports Rate Cap Status Status Status Status Status Status WWPN
----- ---- ---- ------- ------- ------- ------- -------- ------- ------------
1
3Gbps
6Gbps online online online online enabled online 5001086000603824
2
6Gbps

6Gbps offline offline offline offline disabled offline 0000000000000000

```
The following command displays port SFP information about all bridges:
```

cluster1::> system bridge show -sfp
Bridge Name: ATTO_10.226.197.47
Bridge WWN: 2000001086601506
Vendor: Atto
Model: FibreBridge 6500N
Serial Number: FB6500N100526
Firmware Version: 1.60 069G 51.01
Management IP: 10.226.197.47
Errors: -
Last Update Time: 10/22/2015 13:27:37 -04:00

```
FC SFP:
                                                                                Speed
Ports Vendor            Serial Number      Part Number
Capability
----- ----------------- ------------------ ------------------------
----------
    1 AVAGO             AD1020A01FC        AFBR-57D7APZ
8Gbps
    2 AVAGO             AD1020A01F7        AFBR-57D7APZ
8Gbps
```

```
Last Update Timestamp: 10/22/2015 13:27:37 -04:00
```

```
SAS QSFP:
```

```
Ports Vendor            Serial Number       SFP Type       Part Number
----- ----------------- ------------------- --------------
------------------
    1 Molex Inc.        005820292           passive-copper 112-00176
    2 -                 -                   unknown        -
```

```
Last Update Timestamp: -
```

```
Mini-SAS HD:
```

```
Ports Vendor            Serial Number       SFP Type       Part Number
----- ----------------- ------------------- --------------
------------------
    - -                 -                   -              -
```

```
     Bridge Name: ATTO_10.226.197.48
      Bridge WWN: 20000010867002d0
          Vendor: Atto
           Model: FibreBridge 7500N
   Serial Number: FB7500N100018
Firmware Version: 2.00 006U 105.01
   Management IP: 10.226.197.48
          Errors: -
Last Update Time: 10/22/2015 13:27:37 -04:00
```

```
FC SFP:
                                                                    Speed
Ports Vendor            Serial Number      Part Number
Capability
----- ----------------- ------------------ ------------------------
----------
    1 AVAGO             AC1442J00L5        AFBR-57F5MZ
16Gbps
    2 AVAGO             AC1442J00L0        AFBR-57F5MZ
16Gbps
```

```
Last Update Timestamp: -
```

```
SAS QSFP:
```

```
Ports Vendor            Serial Number       SFP Type       Part Number
----- ----------------- ------------------- --------------
------------------
    - -                 -                   -              -
```

```
Last Update Timestamp: 10/22/2015 13:27:37 -04:00
```

```
Mini-SAS HD:
```

```
Ports Vendor            Serial Number       SFP Type       Part Number
----- ----------------- ------------------- --------------
------------------
    1 Amphenol          APF14510026548      Passive Copper 1m ID:00
                                                           112-00429
    2 -                 -                   -              -
    3 -                 -                   -              -
    4 -                 -                   -              -
```

```
The following command displays port statistics information about all
bridges:
```

cluster1::> system bridge show -stats
Bridge Name: ATTO_10.226.197.16
Bridge WWN: 2000001086603824
Vendor: Atto
Model: FibreBridge 6500N
Serial Number: FB6500N101405
Firmware Version: 1.60 A68E 51.01
Management IP: 10.226.197.16
Errors: -
FC Ports:
Oper Neg Link Sync CRC Rx Tx
Ports Status Port Mode Speed Failure Losses Error Words Words
----- ------- ------------- ------- ------- ------ ----- ---------- ----------
1 online n-port 8gb 0 0 0 2721271731 3049186605
2 offline unknown unknown 1 1 0 0 0
Last Update Time: 8/12/2014 12:34:37 -04:00
SAS Ports:
Invalid Disparity Sync PHY Link CRC
SAS PHY Neg Speed Dword Error Loss Reset Changed Error
Port Port Speed Capability Count Count Count Count Count Count
---- ---- ------- ---------- ------- --------- ----- ----- ------- -----
1 0 3Gbps 6Gbps 28262 26665 2 0 1 0
1 1 3Gbps 6Gbps 2110 1794 20 0 1 0
1 2 3Gbps 6Gbps 20435 18857 13 0 1 0
1 3 3Gbps 6Gbps 4573 3353 16 0 1 0
2 0 6Gbps 6Gbps 66 53 0 0 0 0
2 1 6Gbps 6Gbps 27478 25137 2 0 0 0
2 2 6Gbps 6Gbps 20537 17322 9 0 0 0
2 3 6Gbps 6Gbps 23629 21767 10 0 0 0

```
== Related Links

* xref:{relative_path}system-bridge-add.html[system bridge add]




[[ID8d14b501a3b52845b8436e030de27449]]
= system bridge config-dump collect
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Retrieve and save bridge dumpconfiguration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system bridge config-dump collect` command retrieves a
dumpconfiguration file from a system bridge.

== Parameters

``-bridge <text>`` - Bridge Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to retrieve a dumpconfiguration file from the specified
bridge.

== Examples

The following example retrieves a dumpconfiguration file from bridge
ATTO_FibreBridge7500N_1:
```

cluster1::*> system bridge config-dump collect -bridge ATTO_FibreBridge7500N_1
[Job 883] Job is queued: Collect the dumpconfiguration file from bridge "ATTO_FibreBridge7500N_1".

cluster1::*>

```
[[ID906dea6b5f6074d327c0eb6b2ae6cd14]]
= system bridge config-dump delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a dumpconfiguration file

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system bridge config-dump delete` command deletes dumpconfiguration
files previously retrieved with the xref:{relative_path}system-bridge-
config-dump-collect.html[system bridge config-dump collect] command.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to delete a dumpconfiguration file stored on the
specified node.

``-file <text>`` - Config File [.advaced-priv]#(privilege: advanced)#::
Use this parameter to delete the dumpconfiguration file with the specified
file name.

== Examples

The following example deletes dsbridge_config.FB7500N100001.2017-04-
28_14_49_30.txt from node1:
```

cluster1::*> system bridge config-dump delete -node node1 -file dsbridge_config.FB7500N100001.2017-04-28_14_49_30.txt

cluster1::*>

```
== Related Links

* xref:{relative_path}system-bridge-config-dump-collect.html[system bridge
config-dump collect]
```

```
[[IDb33ec1ea67b598936cbd5de288c44f80]]
= system bridge config-dump show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of bridge dumpconfiguration files

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system bridge config-dump show` command displays information about
all the dumpconfiguration files previously retrieved with the
xref:{relative_path}system-bridge-config-dump-collect.html[system bridge
config-dump collect] command. If no parameters are specified, the default
command displays the following information about the dumpconfiguration
files:

* Node
* File Name
* Timestamp
* Bridge
* Bridge Serial Number

To display detailed information about a single dumpconfiguration file, use
the `-node` and `-file` parameters.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Displays information about the dumpconfiguration files stored on the node
```

```
that matches the specified node name.

``[-file <text>]`` - Config File [.advaced-priv]#(privilege: advanced)#::
Displays information about the dumpconfiguration files that match the
specified file name.

``[-bridge <text>]`` - Bridge Name [.advaced-priv]#(privilege:
advanced)#::
Displays information about the dumpconfiguration files from the bridge
that matches the specified bridge name.

``[-serial-number <text>]`` - Serial Number of Bridge [.advaced-
priv]#(privilege: advanced)#::
Displays information about the dumpconfiguration files from the bridge
that matches the specified serial number.

``[-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Time of Collection [.advaced-
priv]#(privilege: advanced)#::
Displays information about the dumpconfiguration files that were collected
at the specified time.

== Examples

The following example displays information about all dumpconfiguration
files:
```

cluster1::*> system bridge config-dump show

Bridge: ATTO_FibreBridge7500N_1

Node File Name Timestamp
----- ---------- ----------
node1
dsbridge_config.FB7500N100001.2017-05-01_09_53_53.txt 5/1/2017 09:53:53
node2
dsbridge_config.FB7500N100001.2017-04-28_14_48_35.txt 4/28/2017 14:48:35
dsbridge_config.FB7500N100001.2017-04-28_15_50_20.txt 4/28/2017 15:50:20
3 entries were displayed.

cluster1::*>

```
The following example displays detailed information about all
dumpconfiguration files:
```

cluster1::*> system bridge config-dump show -instance
Node: node1
Bridge Name: ATTO_FibreBridge7500N_1
Filename: dsbridge_config.FB7500N100001.2017-05-01_09_53_53.txt

```
Timestamp: 5/1/2017 09:53:53
Bridge Serial Number: FB7500N100001
Node: node2
Bridge Name: ATTO_FibreBridge7500N_1
Filename: dsbridge_config.FB7500N100001.2017-04-28_14_48_35.txt
Timestamp: 4/28/2017 14:48:35
Bridge Serial Number: FB7500N100001
Node: node2
Bridge Name: ATTO_FibreBridge7500N_1
Filename: dsbridge_config.FB7500N100001.2017-04-28_15_50_20.txt
Timestamp: 4/28/2017 15:50:20
Bridge Serial Number: FB7500N100001
3 entries were displayed.

cluster1::*>
```

```
== Related Links

* xref:{relative_path}system-bridge-config-dump-collect.html[system bridge
config-dump collect]




[[ID99dc875d839eecb98db20fc0ebf34653]]
= system bridge coredump collect
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Retrieve and save coredump

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system bridge coredump collect` command retrieves a core file from a
bridge.

== Parameters

``-name <text>`` - Bridge Name::
This parameter specifies the bridge name from which the coredump file is
to be collected.

== Examples

The following example retrieves a coredump from bridge
ATTO_FibreBridge7500N_1:
```

cluster1::> system bridge coredump collect -bridge ATTO_FibreBridge7500N_1
[Job 883] Job is queued: Collect the coredump from bridge "ATTO_FibreBridge7500N_1".

cluster1::>

```
[[IDe715c3f0c68f109c44e66752d4f7cf1f]]
= system bridge coredump delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a saved coredump file.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system bridge coredump delete` command deletes a coredump file
previously retrieved with the xref:{relative_path}system-bridge-coredump-
collect.html[system bridge coredump collect] command.

== Parameters

``-name <text>`` - Bridge Name::
This parameter specifies the name of the bridge that the coredump file
belongs to.

``-corename <text>`` - Coredump Filename::
This parameter specifies the name of the coredump file to be deleted.

== Examples

The following example deletes coredump file core.FB7500N100018.1970-01-
05.17_50_30.mem collected from bridge ATTO_FibreBridge7500N_1:
```

cluster1::> system bridge coredump delete -name ATTO_FibreBridge7500N_1 -corename core.FB7500N100018.1970-01-05.17_50_30.mem

cluster1::>

```
== Related Links

* xref:{relative_path}system-bridge-coredump-collect.html[system bridge
coredump collect]
```

```
[[IDa0aac2f5e8c575711b67723905fdcb96]]
= system bridge coredump show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of bridge coredumps

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system bridge coredump show` command displays information about all
the coredump files previously retrieved with the
xref:{relative_path}system-bridge-coredump-collect.html[system bridge
coredump collect] command. If no parameters are specified, the default
command displays the following information about the coredump files:

* Bridge Name
* Bridge Serial Number
* Coredump Filename
* Located on Node
* Panic Timestamp
* Panic String

To display detailed information about a single coredump file, use the `-
node` and `-corename` parameters.

== 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.

``[-serial-number <text>]`` - Bridge Serial Number::
Use this parameter to select the coredump files from the bridge that
matches the specified bridge serial number.

``[-corename <text>]`` - Coredump Filename::
```

```
Use this parameter to select the coredump files that matches the specified
file name.

``[-name <text>]`` - Bridge Name::
Use this parameter to select the coredump files from the bridge that
matches the specified bridge name.

``[-node <nodename>]`` - Located on Node::
Use this parameter to select the coredump the coredump files that are
located on the specified node.

``[-panic-time <MM/DD/YYYY HH:MM:SS>]`` - Panic Timestamp::
Use this parameter to select the coredump files that were collected at the
specified time.

``[-panic-string <text>]`` - Panic String::
Use this parameter to select the coredump files that matches the specified
panic string.

== Examples

The following example displays information about all coredump files:
```

cluster1::> system bridge coredump show
Bridge Name: ATTO_FibreBridge7500N_1
Bridge Serial Number: FB7500N100018
Coredump Filename: core.FB7500N100018.1970-01-05.17_50_30.mem
Located on Node: stg-8020-6a
Panic Timestamp: 7/6/2017 11:03:37
Panic String: CoreDumpGenerate CLI Command

cluster1::>

```
== Related Links

* xref:{relative_path}system-bridge-coredump-collect.html[system bridge
coredump collect]




[[IDfaf745eaf8f030df3486b7d8c807d0a0]]
= system bridge firmware update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Download firmware onto the bridge so it can be updated

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system bridge firmware update` command downloads the firmware onto
the bridge. The bridge needs to be rebooted for the firmware update to
occur. The firmware file to be used is specified by the -uri parameter.

== Parameters

``-bridge <text>`` - Bridge Name [.advaced-priv]#(privilege: advanced)#::
This specifies the bridge whose firmware needs to be updated.

``-uri <text>`` - URI [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the URI from which the firmware file is
downloaded onto the bridge.

``[-skip <true>]`` - Skip Checking for Port Path = 2 [.advaced-
priv]#(privilege: advanced)#::
Use this optional parameter to skip the bridge path checking allowing a
firmware file to be downloaded onto the bridge. Note that doing so might
cause multiple device failures.

``[-username <text>]`` - Username for the URI [.advaced-priv]#(privilege:
advanced)#::
Use this optional parameter to provide a username for the URI.

== Examples

The following example updates the firmware on bridge
ATTO_FibreBridge7500N_1 when credentials are not required.
```

cluster1::> system bridge firmware update -bridge ATTO_FibreBridge7500N_1 -uri http://0.0.0.0/firmware.zbd
[Job 72] Job is queued: System bridge firmware update job.

```
The following example updates the firmware on bridge
ATTO_FibreBridge7500N_1 when credentials are required
```

cluster1::> system bridge firmware update -bridge ATTO_FibreBridge7500N_1 -uri http://0.0.0.0/firmware.zbd
-username user1
Enter the password:

[Job 72] Job is queued: System bridge firmware update job.

```
[[ID476c9b80cba9d3aae55213df1d80c938]]
= system bridge options modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable or disable configurable options for all bridges

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system bridge options modify` command alters the value of
configurable options which affect all the ATTO FibreBridges in the cluster
configuration.

== Parameters

``-option-name <text>`` - Option Name [.advaced-priv]#(privilege:
advanced)#::
The options supported are:
+
  * enable.automatic.addition
+
Allowable values: {true, false}
+
The enable.automatic.addition option controls whether or not ATTO
FibreBridge 7500N and FibreBridge 7600N bridges are automatically added
for in-band monitoring by the Fabric Health Monitor.
+
If set to true, the feature is enabled and all ATTO FibreBridge 7500N and
FibreBridge 7600N bridges in the cluster will automatically be added for
in-band monitoring by the Fabric Health Monitor. This is the preferred and
default value.
+
If set to false, the feature is disabled and ATTO bridges will not
automatically be added for monitoring. This value should only be used if
you do not want to monitor the bridges at all, or if you want them to be
monitored via SNMP.
+
The ATTO FibreBridge 6500N does not have the capability to be monitored by
```

```
in-band management, so this option does not apply to the 6500N.


+
Note that this command is cluster-specific. To affect both clusters of a
MetroCluster system, the command must be executed once on each cluster of
the MetroCluster system.

``[-option-value <text>]`` - Option Value [.advaced-priv]#(privilege:
advanced)#::
This parameter provides the value for each option. Allowable values for
each option are specified in the option description above.

== Examples

The following example sets the enable.automatic.addition option to true:
```

siteA::*> system bridge options modify -option-name enable.automatic.addition -option-value true
siteA::*>

```
[[ID5632f2bde38a5e1b531ead96795cf7a4]]
= system bridge options show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show state of configurable options for all bridges

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The system bridge options show command displays the value of configurable
options which apply to all ATTO FibreBridges in the cluster configuration.
If the user specifies the command without parameters, the output displays
the current value of all the configurable options supported by the
cluster.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-option-name <text>]`` - Option Name [.advaced-priv]#(privilege:
advanced)#::
The options supported are:
+
  * ``_enable.automatic.addition_``
+
Allowable values {true, false}.
+
This option controls whether or not bridges are automatically added for
in-band monitoring by the Fabric Health Monitor.
+
If set to true, the feature is enabled and all ATTO 7500N and 7600N
bridges will automatically be added for in-band monitoring by the Fabric
Health Monitor.
+
If set to false, the feature is disabled and ATTO FibreBridges will not be
automatically added for monitoring. They can be left unmonitored, or
manually added for SNMP monitoring.


+
Note that this command is cluster-specific. To determine the option
setting on both clusters of a MetroCluster system, the command must be
executed once on each cluster of the MetroCluster system.

``[-option-value <text>]`` - Option Value [.advaced-priv]#(privilege:
advanced)#::
The option-value parameter specifies the value of the option. Allowable
values are described for each option supported in the list above.

== Examples

The following example displays the current setting of the bridge options:
```

siteA::*> system bridge options show
Option Name Option Value Option Hint
---------- ---------- ----------
enable.automatic.addition true {true,false} - enable auto-add
siteA::*>

```
:leveloffset: -1


= system chassis commands

:leveloffset: +1


[[IDe992ff476d3dbb60e2e6f3a4734f584c]]
= system chassis show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display all the chassis in the cluster

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system chassis show` command displays information about all the
chassis in the cluster. By default, the command displays the following
information about all the chassis in the cluster:

* Chassis ID
* Status
* List of nodes in the chassis

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about all the chassis in the cluster.

``[-chassis-id <text>]`` - Chassis ID::
Selects information about the specified chassis.

``[-member-nodes {<nodename>|local}]`` - List of Nodes in the Chassis::
Selects information about the chassis with the specified member node list.
```

```
``[-num-nodes <integer>]`` - Number of Nodes in the Chassis::
Selects information about the chassis with the specified number of nodes.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects information about the chassis with the specified status.

== Examples

The following example displays information about all the chassis in the
cluster:
```

cluster1::> system chassis show

Chassis ID Status List of Nodes
--- --- ---
4591227214 ok node1,node2
4591227000 ok node1,node2

```
The following example displays detailed information about a specific
chassis:
```

cluster1::> system chassis show -chassis-id 4591227214 -instance

```
                    Chassis ID: 4591227214
  List of Nodes in the Chassis: node1,node2
Number of Nodes in the Chassis: 2
                        Status: ok
```

```
[[ID3e737254bf651b036702bda3a9c8b8fc]]
= system chassis fru show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the FRUs in the cluster

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `system chassis fru show` command displays information about all the
major chassis specific FRUs in the cluster. By default, the command
displays the following information about all the FRUs in the cluster:

* Chassis ID
* FRU name
* FRU type
* FRU state
* Nodes sharing the FRU

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about FRUs.

``[-node {<nodename>|local}]`` - Node::
Specifies the primary node name in the cluster on which Chassis health
monitor is running.

``[-serial-number <text>]`` - FRU Serial Number::
Selects information about the FRU with the specified serial number.

``[-fru-name <text>]`` - FRU Name::
Selects information about the FRU with the specified FRU name.

``[-type {controller|psu|fan|dimm|bootmedia|ioxm|nvram|nvdimm}]`` - FRU
Type::
Selects information about all the FRUs with the specified FRU type.

``[-name <text>]`` - FRU ID::
Selects information about the FRU with the specified FRU unique name.

``[-state <text>]`` - FRU State::
Selects information about all the FRUs with the specified state.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects information about all the FRUs with the specified status.

``[-display-name <text>]`` - Display Name for the FRU::
Selects information about all the FRUs with the specified FRU display
name.
```

```
``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Monitor Name::
Selects information about all the FRUs with the specified monitor name.

``[-model <text>]`` - Model Type::
Selects information about all the FRUs with the specified FRU model.

``[-shared {shared|not_shared}]`` - Shared Resource::
Selects information about all the FRUs with the specified sharing type.

``[-chassis-id <text>]`` - Chassis ID::
Selects information about all the FRUs in the specified chassis.

``[-additional-info <text>]`` - Additional Information About the FRU::
Selects information about all the FRUs with the specified additional
information.

``[-connected-nodes {<nodename>|local}]`` - List of Nodes Sharing the
FRU::
Selects information about all the FRUs with the specified node list.

``[-num-nodes <integer>]`` - Number of Nodes Sharing the FRU::
Selects information about all the FRUs with the specified number of
connected nodes.

== Examples

The following example displays information about all major chassis
specific FRUs in the cluster:
```

cluster1::> system chassis fru show

Chassis ID FRU Type State Nodes Sharing the FRU
---------- --------- -------- ---------- ----------
4591227214 node1 controller
ok node1
4591227214 node2 controller
ok node2
4591227214 PSU1 FRU psu GOOD node1,node2
4591227214 PSU2 FRU psu GOOD node1,node2

```
The following example displays detailed information about a specific FRU:
```

cluster1::> system chassis fru show -instance -fru-name "PSU1 FRU"
Node: node1

FRU Serial Number: XXT122737891
FRU Name: PSU1 FRU
FRU Type: psu
FRU Name: XXT122737891
FRU State: GOOD
Status: ok
Display Name for the FRU: PSU1 FRU
Monitor Name: chassis
Model Type: none
Shared Resource: shared
Chassis ID: 4591227214
Additional Information About the FRU: Part Number: 114-00065+A0
Revision: 020F
Manufacturer: NetApp
FRU Name: PSU
List of Nodes Sharing the FRU: node1,node2
Number of Nodes Sharing the FRU: 2

```
:leveloffset: -1


= system configuration commands

:leveloffset: +1


[[ID9a48f98170bc424a25607b3edf8a32ca]]
= system configuration backup copy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Copy a configuration backup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup copy` command copies a configuration
backup from one node in the cluster to another node in the cluster.

Use the xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show] command to display configuration backups to
copy.

== Parameters
```

```
``-from-node {<nodename>|local}`` - Source Node [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify the name of the source node where the
configuration backup currently exists.

``-backup <text>`` - Backup Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the name of the configuration backup file to
copy.

``-to-node {<nodename>|local}`` - Destination Node [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify the name of the destination node where the
configuration backup copy is created.

== Examples

The following example copies the configuration backup file
`node1.special.7z` from the node `node1` to the node `node2` .
```

cluster1::*> system configuration backup copy -from-node node1 -backup node1.special.7z -to-node node2
[Job 295] Job is queued: Copy backup job.

```
== Related Links

* xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show]




[[ID3fb57c0b535a6c2cea2f7c197ee11f04]]
= system configuration backup create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a configuration backup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup create` command creates a new
configuration backup file.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node on which to create the backup file.

``[-backup-name <text>]`` - Backup Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the name of the backup file to create. The
backup name cannot contain a space or any of the following characters: * ?
/

``[-backup-type {node|cluster}]`` - Backup Type [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify the type of backup file to create.

== Examples

The following example creates a a new cluster configuration backup file
called `node1.special.7z` on the node `node1` .
```

cluster1::*> system configuration backup create -node node1 -backup-name node1.special.7z -backup-type cluster
[Job 194] Job is queued: Cluster Backup OnDemand Job.

```
[[IDe97b4916d976497e372497f28f35c99a]]
= system configuration backup delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a configuration backup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup delete` command deletes a saved
configuration backup.

Use the xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show] command to display saved configuration backups.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the name of the source node where the
configuration backup currently exists.

``-backup <text>`` - Backup Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the name of the configuration backup file to
delete.

== Examples

The following example shows how to delete the configuration backup file
`node1.special.7z` from the node `node1` .
```

cluster1::*> system configuration backup delete -node node1 -backup node1.special.7z

```
== Related Links

* xref:{relative_path}system-configuration-backup-show.html[system
```

```
configuration backup show]




[[ID25ede7786db5452c9a0af07b23619e4d]]
= system configuration backup download
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Download a configuration backup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup download` command copies a configuration
backup from a source URL to a node in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the name of the node to which the
configuration backup is downloaded.

``-source <text>`` - Source URL [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the source URL of the configuration backup
to download.

``[-backup-name <text>]`` - Backup Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify a new local file name for the downloaded
configuration backup.

``[-validate-certificate {true|false}]`` - Validate Digital Certificate
[.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value ``_true_`` to validate the digital
certificate of the remote server.

== Examples

The following example shows how to download a configuration backup file
```

```
from a URL to a file named `exampleconfig.download.7z` on the node `node2`
.
```

cluster1::*> system configuration backup download -node node2 -source http://www.example.com/config/download/nodeconfig.7z -backup-name exampleconfig.download.7z

```
[[ID33b70e34464d677d40a90e860e2c3641]]
= system configuration backup rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a configuration backup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup rename` command changes the file name of
a configuration backup file.

Use the xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show] command to display configuration backups to
rename.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the name of the source node where the
configuration backup currently exists.

``-backup <text>`` - Backup Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the name of the configuration backup file to
rename.

``-new-name <text>`` - New Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a new name for the configuration backup
file.

== Examples

The following example renames the saved configuration file
`download.config.7z` on the node `node1` to `test.config.7z` .
```

cluster1::*> system configuration backup rename -node node1 -backup download.config.7z -new-name test.config.7z

```
== Related Links

* xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show]




[[IDab02367185d8d4decbd3f9eb896db1af]]
= system configuration backup show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show configuration backup information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup show` command displays information about
saved configuration backups.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects configuration backups that are saved on the node you specify.

``[-backup <text>]`` - Backup Name [.advaced-priv]#(privilege:
advanced)#::
Selects configuration backups that have the backup name you specify.

``[-backup-type {node|cluster}]`` - Backup Type [.advaced-
priv]#(privilege: advanced)#::
```

```
Selects configuration backups of the type you specify.

``[-time <MM/DD HH:MM:SS>]`` - Backup Creation Time [.advaced-
priv]#(privilege: advanced)#::
Selects configuration backups that were saved on the date and time you
specify.

``[-cluster-name <text>]`` - Cluster Name [.advaced-priv]#(privilege:
advanced)#::
Selects configuration backups that were saved in the cluster that has the
name you specify.

``[-cluster-uuid <UUID>]`` - Cluster UUID [.advaced-priv]#(privilege:
advanced)#::
Selects configuration backups that were saved in the cluster that has the
UUID you specify.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of Backup [.advaced-
priv]#(privilege: advanced)#::
Selects configuration backups that have the file size you specify.

``[-nodes-in-backup {<nodename>|local}]`` - Nodes In Backup [.advaced-
priv]#(privilege: advanced)#::
Selects configuration backups that include the configuration of the nodes
you specify.

``[-version <text>]`` - Software Version [.advaced-priv]#(privilege:
advanced)#::
Selects configuration backups that have the software version you specify.

``[-is-auto {true|false}]`` - Backup Created from Schedule (true or false)
[.advaced-priv]#(privilege: advanced)#::
A value of true selects configuration backups that were created from a
schedule. A value of false selects configuration backups that were created
manually.

``[-schedule <text>]`` - Name of Backup Schedule [.advaced-
priv]#(privilege: advanced)#::
Selects configuration backups that were created by the schedule you
specify.

== Examples

The following example shows typical output for this command.
```

cluster1::*> system configuration backup show

Node Backup Tarball Time Size
--------- ---------- ---------- -----
node1 cluster1.8hour.2011-02-22.18_15_00.7z 02/22 18:15:00 7.78MB
node1 cluster1.8hour.2011-02-23.02_15_00.7z 02/23 02:15:00 7.98MB
node1 cluster1.8hour.2011-02-23.10_15_00.7z 02/23 10:15:00 7.72MB
node1 cluster1.daily.2011-02-22.00_10_00.7z 02/22 00:10:00 7.19MB
node1 cluster1.daily.2011-02-23.00_10_00.7z 02/23 00:10:00 7.99MB
Press <space> to page down, <return> for next line, or 'q' to quit… q
5 entries were displayed.

```
[[IDa75e839848b61000e12cc8aabeb6dabf]]
= system configuration backup upload
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Upload a configuration backup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup upload` command copies a configuration
backup from a node in the cluster to a remote URL.

NOTE: the web server to which you are uploading the configuration backup
file must have PUT operations enabled for HTTP and POST operations enabled
for HTTPS. Some web servers may require the installation of an additional
module. For more information, see your web server's documentation.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the name of the node from which the
configuration backup is uploaded.

``-backup <text>`` - Backup Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the file name of the configuration backup to
upload.

``-destination <text>`` - Destination URL [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the destination URL of the configuration
```

```
backup.

``[-validate-certificate {true|false}]`` - Validate Digital Certificate
[.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value ``_true_`` to validate the digital
certificate of the remote server.

``[-rest-method {POST|PUT}]`` - REST Method [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the http method to use in order to upload
the configuration backup.

== Examples

The following example show how to upload the configuration backup file
`testconfig.7z` from the node `node2` to a remote URL.
```

cluster1::*> system configuration backup upload -node node2 -backup testconfig.7z -destination ftp://www.example.com/config/uploads/testconfig.7z

```
[[ID5d9df7a382b0585e4758c9df82def854]]
= system configuration backup settings clear-password
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear password for destination URL

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup settings clear-password ` command clears
the password used for uploads of configuration backups. Enter the command
without parameters.

Use the xref:{relative_path}system-configuration-backup-settings-
show.html[system configuration backup settings show] command to display
the destination URL for configuration backups. Use the
xref:{relative_path}system-configuration-backup-settings-
modify.html[system configuration backup settings modify] command to change
the destination URL and remote username for configuration backups.

== Parameters

== Examples

The following example shows successful execution of this command.
```

cluster1::*> system configuration backup settings clear-password

Do you want to clear the password? {y|n}:

```
== Related Links

* xref:{relative_path}system-configuration-backup-settings-
show.html[system configuration backup settings show]
* xref:{relative_path}system-configuration-backup-settings-
modify.html[system configuration backup settings modify]
```

```
[[IDd018658f02b2f336b477a485e2030afc]]
= system configuration backup settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify configuration backup settings

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup settings modify` command changes settings
for configuration backup.

== Parameters

``[-destination <text>]`` - Backup Destination URL [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify the destination URL for uploads of
configuration backups. Use the value "" to remove the destination URL.
Only FTP, HTTP, HTTPS, TFTP, FTPS protocols are supported.

``[-username <text>]`` - Username for Destination [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify the user name to use to log in to the
destination system and perform the upload. Use the
xref:{relative_path}system-configuration-backup-settings-set-
password.html[system configuration backup settings set-password] command
to change the password used with this user name.

``[-validate-certificate {true|false}]`` - Validate Digital Certificate
[.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value ``_true_`` to validate the digital
certificate of the remote server. Digital certificate validation is
disabled by default.

``[-numbackups1 <integer>]`` - Number of Backups to Keep for Schedule 1
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the number of backups created by backup job
schedule 1 to keep on the destination system. If the number of backups
exceeds this number, the oldest backup is removed. Schedule 1 is pre-
```

```
programmed to be '8hour'.

``[-numbackups2 <integer>]`` - Number of Backups to Keep for Schedule 2
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the number of backups created by backup job
schedule 2 to keep on the destination system. If the number of backups
exceeds this number, the oldest backup is removed. Schedule 2 is pre-
programmed to be 'daily'.

``[-numbackups3 <integer>]`` - Number of Backups to Keep for Schedule 3
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the number of backups created by backup job
schedule 3 to keep on the destination system. If the number of backups
exceeds this number, the oldest backup is removed. Schedule 3 is pre-
programmed to be 'weekly'.

== Examples

The following example shows how to set the destination URL and user name
used for uploads of configuration backups.
```

cluster1::*> system configuration backup settings modify -destination ftp://www.example.com/config/uploads/ -username admin

```
== Related Links

* xref:{relative_path}system-configuration-backup-settings-set-
password.html[system configuration backup settings set-password]




[[IDcf012c67e7f2717f7932f52ad84a250f]]
= system configuration backup settings set-password
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify password for destination URL

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup settings set-password` command sets the
password used for uploads of configuration backups. This password is used
along with the username you specify using the xref:{relative_path}system-
configuration-backup-settings-modify.html[system configuration backup
settings modify] command to log in to the system and perform the upload.
Enter the command without parameters. The command prompts you for a
password, and for a confirmation of that password. Enter the same password
at both prompts. The password is not displayed.

Use the xref:{relative_path}system-configuration-backup-settings-
show.html[system configuration backup settings show] command to display
the destination URL for configuration backups. Use the
xref:{relative_path}system-configuration-backup-settings-
modify.html[system configuration backup settings modify] command to change
the destination URL and remote username for configuration backups.

== Parameters

== Examples

The following example shows successful execution of this command.
```

cluster1::*> system configuration backup settings set-password

Enter the password:
Confirm the password:

```
== Related Links

* xref:{relative_path}system-configuration-backup-settings-
modify.html[system configuration backup settings modify]
* xref:{relative_path}system-configuration-backup-settings-
show.html[system configuration backup settings show]




[[ID1d6b84b2aab16c27f1bb5f54b55f890b]]
= system configuration backup settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show configuration backup settings

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration backup settings show` command displays current
settings for configuration backup. These settings apply to backups created
automatically by schedules. By default, the command displays the URL to
which configuration backups are uploaded, and the user name on the remote
system used to perform the upload.

Use the xref:{relative_path}system-configuration-backup-settings-set-
password.html[system configuration backup settings set-password] command
to change the password used with the user name on the destination. Use the
xref:{relative_path}system-configuration-backup-settings-
modify.html[system configuration backup settings modify] command to change
the destination URL and username for uploads of configuration backups, and
to change the number of backups to keep for each schedule.

== Parameters

``[-instance  ]`` [.advaced-priv]#(privilege: advanced)#::
```

```
Use this parameter to display detailed information about configuration
backup settings, including the number of backups to keep for each backup
schedule.

== Examples

The following example displays basic backup settings information.
```

cluster1::*> system configuration backup settings show
Backup Destination URL Username
--- ---
ftp://www.example.com/config/uploads/ jdoe

```
The following example shows detailed output using the ``_-instance_``
parameter.
```

cluster1::*> system configuration backup settings show -instance
Backup Destination URL: ftp://www.example.com/config/uploads/
Username for Destination: admin
Validate Digital Certificate: true
Schedule 1: 8hour
Number of Backups to Keep for Schedule 1: 2
Schedule 2: daily
Number of Backups to Keep for Schedule 2: 2
Schedule 3: weekly
Number of Backups to Keep for Schedule 3: 2

```
== Related Links

* xref:{relative_path}system-configuration-backup-settings-set-
password.html[system configuration backup settings set-password]
* xref:{relative_path}system-configuration-backup-settings-
modify.html[system configuration backup settings modify]




[[ID4890178f61d93814beb2aa3b2f73a301]]
= system configuration recovery cluster modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify cluster recovery status

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration recovery cluster modify` command modifies the
cluster recovery status. This command should be used to end the cluster
recovery after all recovery procedures are applied.

== Parameters

``[-recovery-status {complete|in-progress|not-in-recovery}]`` - Cluster
Recovery Status [.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value `complete` to set the cluster recovery
status after the cluster has been recreated and all of the nodes have been
rejoined to it. This enables each node to resume normal system operations.
The `in-progress` and `not-in-recovery` values are not applicable to this
command.

== Examples

The following example modifies the cluster recovery status.
```

source::> system configuration recovery cluster modify -recovery-status complete

```
[[IDaf126b8016699659f9d9466a7fcce77a]]
= system configuration recovery cluster recreate
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Recreate cluster

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration recovery cluster recreate` command re-creates a
cluster, using either the current node or a configuration backup as a
configuration template. After you re-create the cluster, rejoin nodes to
the cluster using the xref:{relative_path}system-configuration-recovery-
cluster-rejoin.html[system configuration recovery cluster rejoin] command.

== Parameters

``-from {node|backup}`` - From Node or Backup [.advaced-priv]#(privilege:
advanced)#::
Use this parameter with the value `node` to re-create the cluster using
the current node as a configuration template. Use this parameter with the
value `backup` to re-create the cluster using a configuration backup as a
configuration template.

``[-backup <text>]`` - Backup Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the name of a configuration backup file to
use as a configuration template. If you specified the ``_-from_``
parameter with the value `backup` , you must use this parameter and
specify a backup name. Use the xref:{relative_path}system-configuration-
backup-show.html[system configuration backup show] command to view
available configuration backup files.

== Examples

The following example shows how to re-create a cluster using the node
`node1` as a configuration template.
```

cluster1::*> system configuration recovery cluster recreate -from node

```
The following example shows how to re-create a cluster using the
configuration backup `siteconfig.backup.7z` as a configuration template.
```

cluster1::*> system configuration recovery cluster recreate -from backup -backup siteconfig.backup.7z

```
== Related Links

* xref:{relative_path}system-configuration-recovery-cluster-
rejoin.html[system configuration recovery cluster rejoin]
* xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show]




[[IDb1c6683d4fb88164f6d311608f53d5ea]]
= system configuration recovery cluster rejoin
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rejoin a cluster

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration recovery cluster rejoin` command rejoins a node
to a new cluster created earlier using the xref:{relative_path}system-
configuration-recovery-cluster-recreate.html[system configuration recovery
cluster recreate] command. Only use this command to recover a node from a
disaster. Because this synchronization can overwrite critical cluster
information, and will restart the node you specify, you are required to
confirm this command before it executes.

== Parameters

``-node {<nodename>|local}`` - Node to Rejoin [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node to rejoin to the cluster.

== Examples

This example shows how to rejoin the node `node2` to the cluster.
```

cluster1::*> system configuration recovery cluster rejoin -node node2

Warning: This command will rejoin node "node2" into the local cluster,

potentially overwriting critical cluster configuration files. This
command should only be used to recover from a disaster. Do not perform
any other recovery operations while this operation is in progress.
This command will cause node "node2" to reboot.
Do you want to continue? {y|n}: y

```
== Related Links

* xref:{relative_path}system-configuration-recovery-cluster-
recreate.html[system configuration recovery cluster recreate]




[[ID36b63b79e9f203bac7425f1269b5d23a]]
= system configuration recovery cluster show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show cluster recovery status

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration recovery cluster show` command displays the
cluster recovery status. Cluster recovery status is "not-in-recovery"
under normal operations, and it becomes "in-progress" if a new cluster is
created using the xref:{relative_path}system-configuration-recovery-
cluster-recreate.html[system configuration recovery cluster recreate]
command with the ``_-from backup_`` parameter. When cluster recovery
status is "in-progress", wait until the output of the "Is Recovery Status
Persisted" field is true before using the xref:{relative_path}system-
configuration-recovery-cluster-rejoin.html[system configuration recovery
cluster rejoin] command to recover other nodes in the cluster.

== Examples

The following example displays the cluster recovery status.
```

source::> system configuration recovery cluster show
Recovery Status: in-progress
Is Recovery Status Persisted: true

```
== Related Links

* xref:{relative_path}system-configuration-recovery-cluster-
recreate.html[system configuration recovery cluster recreate]
* xref:{relative_path}system-configuration-recovery-cluster-
rejoin.html[system configuration recovery cluster rejoin]




[[IDbdb38f7d587ea7640026d0cee924a1cf]]
= system configuration recovery cluster sync
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Sync a node with cluster configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration recovery cluster sync` command synchronizes a
node with the cluster configuration. Only use this command to recover a
node from a disaster. Because this synchronization can overwrite critical
cluster information, and will restart the node you specify, you are
required to confirm this command before it executes.

== Parameters

``-node {<nodename>|local}`` - Node to Synchronize [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify the name of the node to synchronize with the
cluster.

== Examples

The following example shows the synchronization of the node `node2` to the
cluster configuration.
```

cluster1::*> system configuration recovery cluster sync -node node2

Warning: This command will synchronize node "node2" with the cluster

configuration, potentially overwriting critical cluster configuration
files on the node. This feature should only be used to recover from a
disaster. Do not perform any other recovery operations while this
operation is in progress. This command will cause all the cluster
applications on node "node2" to restart, interrupting administrative
CLI and Web interface on that node.
Do you want to continue? {y|n}: y
All cluster applications on node "node2" will be restarted. Verify that the cluster applications go online.

```
[[ID0d08e1c7b1c8715eeeee1e8d0398f7d3]]
= system configuration recovery node restore
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore node configuration from a backup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system configuration recovery node restore` command restores the
configuration of the local node from a configuration backup file.

Use the xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show] command to view available configuration backup
files.

== Parameters

``-backup <text>`` - Backup Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the name of a configuration backup file to
use as the configuration template.

``[-nodename-in-backup <text>]`` - Use Backup Identified by this Nodename
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a node within the configuration backup file
to use as a configuration template. Only specify this parameter if you are
specifying a name other than the name of the local node.

``[-force <true>]`` - Force Restore Operation [.advaced-priv]#(privilege:
advanced)#::
Use this parameter with the value `true` to force the restore operation
and overwrite the current configuration of the local node. This overrides
```

```
all compatibility checks between the node and the configuration backup.
The configuration in the backup is installed even if it is not compatible
with the node's software and hardware.
+
Use this parameter with the value `false` to be warned of the specific
dangers of restoring and be prompted for confirmation before executing the
command. This value also assures that the command performs compatibility
checks between configuration stored in the backup and the software and
hardware of the node. The default is `false` .

== Examples

The following example shows how to restore the configuration of the local
node from the configuration backup of `node3` that is stored in the
configuration backup file `example.backup.7z` .
```

cluster1::*> system configuration recovery node restore -backup example.backup.7z
Warning: This command overwrites local configuration files with files contained
in the specified backup file. Use this command only to recover from a
disaster that resulted in the loss of the local configuration files.
The node will reboot after restoring the local configuration.
Do you want to continue? {y|n}: y

```
== Related Links

* xref:{relative_path}system-configuration-backup-show.html[system
configuration backup show]




:leveloffset: -1


= system controller commands

:leveloffset: +1


[[ID8b27054183ad90dac99aefeca734c6cc]]
= system controller show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display the controller information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller show` command displays information about all the
controllers in the cluster. These commands are available for 80xx, 25xx
and later systems. Earlier models are not supported. By default, the
command displays the following information about all the controllers in
the cluster:

* Controller name
* System ID
* System serial number
* Controller model name
* Health monitor status

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about all the controllers in the cluster.

``[-node {<nodename>|local}]`` - Node::
Selects information about the specified controller.

``[-system-id <text>]`` - System ID::
Selects information about the controller with the specified System ID.

``[-model <text>]`` - Model Name::
Selects information about the controllers with the specified model name.

``[-part-number <text>]`` - Part Number::
Selects information about the controllers with the specified part number.

``[-revision <text>]`` - Revision::
Selects information about the controllers with the specified revision.

``[-serial-number <text>]`` - Serial Number::
Selects information about the controller with the specified system serial
```

```
number.

``[-controller-type <text>]`` - Controller Type::
Selects information about the controllers with the specified controller
type.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects information about the controllers with the specified health
monitor status.

``[-chassis-id <text>]`` - Chassis ID::
Selects information about the controllers with the specified chassis ID.

== Examples

The below example displays information about all controllers in the
cluster.
```

cluster1::> system controller show
Controller Name System ID Serial Number Model Status
---------- ---------- ---------- -------- ----------
node1 140733730268652 700001456939 FAS2520 ok
node2 140733730268667 700001456941 FAS2520 ok
2 entries were displayed.

```
The example below displays detailed information about specified controller
in the cluster.
```

cluster1::> system controller show -instance -node node1
Node: node1
System ID: 140733730268652
Model Name: FAS2520
Part Number: 111-01316
Revision: 21
Serial Number: 700001456939
Controller Type: none
Status: ok
Chassis ID: 4591227214

```
[[ID33668b2305ed93d111a9569b8b653d4d]]
= system controller bezel-led modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the state of the bezel LED log (lit or unlit)

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller bezel-led modify` command controls the illumination
state of the LED logo on the chassis bezel (lit or unlit) of a node. On
platforms where two controllers in a high-availability (HA) configuration
share the same chassis, any changes to the LED logo's illumination state
are kept synchronized on both nodes.

== Parameters

``-node {<nodename>|local}`` - Node::
The node whose bezel LED logo illumination state is being modified.

``[-state {off|on|disconnected|bezel-switch|unknown}]`` - Bezel LED Logo
State::
The bezel LED logo illumination state of the node. The possible values
are: on (lit), off (unlit),bezel-switch or disconnected. The default state
is `on` .

== Examples

The following example turns on the bezel LED logo:
```

cluster1::> system controller bezel-led modify -node node1 -state on

```
The following example turns off the bezel LED logo:
```

cluster1::> system controller bezel-led modify -node node1 -state off

```
The following example lets the bezel switch control the bezel LED:
```

cluster1::> system controller bezel-led modify -node node1 -state bezel-switch

```
[[ID864e0e0b55db4b13a1dde0ee699b0f9d]]
= system controller bezel-led show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the state of the bezel LED logo (lit or unlit)

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller bezel-led show` command shows the current
illumination state of the bezel LED logo (lit or unlit).

== 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>|local}]`` - Node::
The node whose bezel LED logo illumination state is being shown.

``[-state {off|on|disconnected|bezel-switch|unknown}]`` - Bezel LED Logo
State::
The bezel LED logo illumination state of the node.

== Examples

The following example lists the current illumination state of the bezel
LED logo:
```

cluster1::> system controller bezel-led show
Node Bezel LED Logo State
---------- ----------

node1 off
node2 off

```
[[ID23ea24b22eaa0e9a9e4b69228b81e6c9]]
= system controller bootmedia show-percentage-used
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the utilization percentage of the NVMe boot media

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller bootmedia show-percentage-used` command displays
the Boot Media Device percentage used. These commands are available for
NVMe bootmedia systems. By default, the command displays the following
information about the NVMe bootmedia:

* Node
* Display Name
* Is NVMe
* NVMe Percentage Used
* NVMe Available Spare
* NVMe Spare Threshold

To display more details, use the `-instance` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for all the boot media devices.

``[-node {<nodename>|local}]`` - Node::
Selects the bootmedia device that is present on the specified node.

``[-serial-num <text>]`` - Serial Number::
Selects the bootmedia devices with the specified serial number.
```

```
``[-vendor-id <Hex Integer>]`` - Vendor ID::
Selects the bootmedia devices with the specified vendor ID.

``[-device-id <Hex Integer>]`` - Device ID::
Selects the bootmedia devices with the specified device ID.

``[-display-name <text>]`` - Display Name::
Selects the bootmedia devices with the specified display name.

``[-unique-name <text>]`` - Unique Name::
Selects the bootmedia device with the specified unique name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the bootmedia devices with the specified health monitor.

``[-usbmon-status {present|not-present}]`` - Bootmedia Health Monitor::
Selects the bootmedia devices with the specified USBMON status.

``[-device-state {good|warn|bad}]`` - Bootmedia State::
Selects the bootmedia devices with the specified device state.

``[-size <integer>]`` - Max Memory Size (MB)::
Selects the bootmedia devices with the specified memory size.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the bootmedia devices with the specified health monitor status.

``[-is-nvme {true|false}]`` - Is NVMe Boot Media::
Selects the boot media technology. If the value is "true", the boot media
uses NVMe technology.

``[-nvme-percent-used <integer>]`` - NVMe Percentage Used::
Selects the NVMe boot media with the specified percentage of rated life
used by the solid-state device (SSD). When an SSD reaches 100% of rated
life, it might not be able to retain data while powered off for long
periods of time.

``[-nvme-available-spare <integer>]`` - NVMe Available Spare::
Selects the NVMe boot media with the specified percentage of available
spare blocks. When all spare blocks are consumed on the SSD, i.e. when
this parameter reaches 0%, the system can no longer write data to the
solid-state device (SSD).

``[-nvme-spare-threshold <integer>]`` - NVMe Spare Threshold::
```

```
Selects the NVMe boot media with the specified "low spare blocks"
threshold percentage. System health monitor alerts are generated when the
available spare blocks on the SSD cross below this threshold.

== Examples

The following example displays the information of the boot media devices
present in all the nodes in a cluster:
```

cluster1::> system controller bootmedia show-percentage-used

+ NVMe Used NVMe Spare NVMe Spare
+Node Display Name Is NVMe Percentage Available Threshold
+---------- -------- ------- ---------- ---------- ----------
localhost
BootMedia/ true 0 100 10
0X331511900503A0SAM0
00PM9A30002T00025000

```
The following example displays the detailed information about the boot
media present in a node:
```

cluster1::> system controller bootmedia show-percentage-used -instance -node node1
Node: localhost
Serial Number: S668NE0RC00442
Vendor ID: 144d
Device ID: a80a
Display Name: BootMedia/0X331511900503A0SAM000PM9A30002T00025000
Unique Name: 0X331511900503A0SAM000PM9A30002T00025000
Health Monitor Name: controller
Bootmedia Health Monitor: present
Bootmedia State: good
Max Memory Size (MB): 1831420
Status: ok
Is NVME Boot Media: true
NVME Percentage Used: 0
NVME Available Spare: 100
NVME Spare Threshold: 10

```
[[ID81697fd3c2a62360353f729de1b10872]]
= system controller bootmedia show-serial-number
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Boot Media Device serial number
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller bootmedia show-serial-number` command displays the
Boot Media Device serial number. These commands are available for 80xx,
25xx and later systems. Earlier models are not supported. By default, the
command displays the following information about the bootmedia:

* Node
* Display Name
* Serial Number
* Size
* Bootmedia State
* Status

To display more details, use the `-instance` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for all the bootmedia devices.

``[-node {<nodename>|local}]`` - Node::
Selects the bootmedia device that is present on the specified node.

``[-serial-num <text>]`` - Serial Number::
Selects the bootmedia devices with the specified serial number.

``[-vendor-id <Hex Integer>]`` - Vendor ID::
Selects the bootmedia devices with the specified vendor ID.

``[-device-id <Hex Integer>]`` - Device ID::
Selects the bootmedia devices with the specified device ID.

``[-display-name <text>]`` - Display Name::
Selects the bootmedia devices with the specified display name.

``[-unique-name <text>]`` - Unique Name::
Selects the bootmedia device with the specified unique name.

``[-monitor {node-connect|system-
```

```
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the bootmedia devices with the specified health monitor.

``[-usbmon-status {present|not-present}]`` - Bootmedia Health Monitor::
Selects the bootmedia devices with the specified USBMON status.

``[-device-state {good|warn|bad}]`` - Bootmedia State::
Selects the bootmedia devices with the specified device state.

``[-size <integer>]`` - Max Memory Size (MB)::
Selects the bootmedia devices with the specified memory size.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the bootmedia devices with the specified health monitor status.

``[-is-nvme {true|false}]`` - Is NVMe Boot Media::
Selects the boot media technology. If the value is "true", the boot media
uses NVMe technology.

``[-nvme-percent-used <integer>]`` - NVMe Percentage Used::
Selects the NVMe boot media with the specified percentage of rated life
used by the solid-state device (SSD). When an SSD reaches 100% of rated
life, it might not be able to retain data while powered off for long
periods of time.

``[-nvme-available-spare <integer>]`` - NVMe Available Spare::
Selects the NVMe boot media with the specified percentage of available
spare blocks. When all spare blocks are consumed on the SSD, i.e. when
this parameter reaches 0%, the system can no longer write data to the
solid-state device (SSD).

``[-nvme-spare-threshold <integer>]`` - NVMe Spare Threshold::
Selects the NVMe boot media with the specified "low spare blocks"
threshold percentage. System health monitor alerts are generated when the
available spare blocks on the SSD cross below this threshold.

== Examples

The following example displays the information of the boot media present
in all the nodes in a cluster:
```

cluster1::> system controller bootmedia show-serial-number

Node Display Name Serial Number (MB) State Status
---------- ---------- ---------- ------- --------- ------

or-12-01
BootMedia/SAMSUNG S2J4NXAGA08186 122104 good ok
MZVLV128HCGR-00000
BootMedia-2/SAMSUNG S2J4NXAGA08198 122104 good ok
MZVLV128HCGR-00000
2 entries were displayed.

```
The following example displays the detailed information about the boot
media present in a node:
```

cluster1::> system controller bootmedia show-serial-number -instance -node node1
Node: node1
Vendor ID: 8086
Device ID: 8d02
Display Name: TOSHIBA THNSNJ060GMCU
Unique Name: /dev/ad4s1 (TOSHIBA THNSNJ060GMCU)
Health Monitor Name: controller
Bootmedia Health Monitor: present
Bootmedia State: good
Max memory size(in MB): 16367
Status: ok
Serial number: Y4IS104FTNEW
Is NVMe: true
NVMe Percentage Used: 0
NVMe Available Spare: 100
NVMe Spare Threshole: 10

```
[[ID14117ca5dd99f0437ae4c41f64f2cac6]]
= system controller bootmedia show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Boot Media Device Health Status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller bootmedia show` command displays details of the
bootmedia devices present in all the nodes in a cluster. These commands
are available for 80xx, 25xx and later systems. Earlier models are not
supported. By default, the command displays the following information
about the bootmedia:
```

```
* Node name
* Display name
* Vendor ID
* Device ID
* Memory size
* Bootmedia state
* Health monitor status

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for all the bootmedia devices.

``[-node {<nodename>|local}]`` - Node::
Selects the bootmedia device that is present on the specified node.

``[-serial-num <text>]`` - Serial Number::
Selects the bootmedia devices with the specified serial number.

``[-vendor-id <Hex Integer>]`` - Vendor ID::
Selects the bootmedia devices with the specified vendor ID.

``[-device-id <Hex Integer>]`` - Device ID::
Selects the bootmedia devices with the specified device ID.

``[-display-name <text>]`` - Display Name::
Selects the bootmedia devices with the specified display name.

``[-unique-name <text>]`` - Unique Name::
Selects the bootmedia device with the specified unique name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the bootmedia devices with the specified health monitor.

``[-usbmon-status {present|not-present}]`` - Bootmedia Health Monitor::
Selects the bootmedia devices with the specified USBMON status.

``[-device-state {good|warn|bad}]`` - Bootmedia State::
Selects the bootmedia devices with the specified device state.
```

```
``[-size <integer>]`` - Max Memory Size (MB)::
Selects the bootmedia devices with the specified memory size.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the bootmedia devices with the specified health monitor status.

``[-is-nvme {true|false}]`` - Is NVMe Boot Media::
Selects the boot media technology. If the value is "true", the boot media
uses NVMe technology.

``[-nvme-percent-used <integer>]`` - NVMe Percentage Used::
Selects the NVMe boot media with the specified percentage of rated life
used by the solid-state device (SSD). When an SSD reaches 100% of rated
life, it might not be able to retain data while powered off for long
periods of time.

``[-nvme-available-spare <integer>]`` - NVMe Available Spare::
Selects the NVMe boot media with the specified percentage of available
spare blocks. When all spare blocks are consumed on the SSD, i.e. when
this parameter reaches 0%, the system can no longer write data to the
solid-state device (SSD).

``[-nvme-spare-threshold <integer>]`` - NVMe Spare Threshold::
Selects the NVMe boot media with the specified "low spare blocks"
threshold percentage. System health monitor alerts are generated when the
available spare blocks on the SSD cross below this threshold.

== Examples

The following example displays the information of the boot media present
in all the nodes in a cluster:
```

cluster1::> system controller bootmedia show
Size Bootmedia
Node Display Name Vendor ID Device ID (MB) State Status
---------- ---------- --------- --------- ------- --------- ------
node1 Micron Technology 634 655 1929 good ok
0x655
node2 Micron Technology 634 655 1929 good ok
0x655

```
The example below displays the detailed information about the bootmedia
present in a node.
```

cluster1::> system controller bootmedia show -instance -node node1

Node: node1
Vendor ID: 634
Device ID: 655
Display Name: Micron Technology 0x655
Unique Name: Micron Technology 0x655 (ad.0)
Health Monitor Name: controller
USBMON Health Monitor: present
Bootmedia State: good
Max memory size(in MB): 1929
Status: ok
Is NVMe: true
NVMe Percentage Used: 0
NVMe Available Spare: 100
NVMe Spare Threshold: 10

```
[[IDaaa05538d2287a4e69f1ea9faeb7013f]]
= system controller clus-flap-threshold show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the controller cluster port flap threshold

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller clus-flap-threshold show` command allows the
display of the threshold for link flapping counts for all nodes. This
threshold would be the number of times the cluster port links for a given
node can flap (go down) within a polling period before triggering an
alert.




[[IDffc2fdbebc3f1ecafcddf65691bc0a2a]]
= system controller config show-errors
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display configuration errors
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller config show-errors` displays configuration errors.

* node
* description

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for all the PCI devices.

``[-node {<nodename>|local}]`` - Node::
Displays configuration errors on the specified node.

``[-verbose <true>]`` - Verbose Output?::
The `-verbose` parameter enables verbose mode, resulting in the display of
more detailed output.

``[-description <text>]`` - Error Description::
Displays the node with the specified configuration error.

== Examples

The example below displays configuration errors on all the nodes in the
cluster.
```

cluster1::> system controller config show-errors

Configuration Info and Errors for Node: cluster1-01

Chelsio T320E 2x10G NIC card (PN X1008A) in slot 1 is not supported on model FAS3210

Configuration Info and Errors for Node: cluster1-02

PCI-E Dual 10/100/1000 Ethernet G20 card (PN X1039A) in slot 2 is not supported on model FAS3210

cluster1::>

```
cluster1::> system controller config show-errors -verbose

Configuration Info and Errors for Node: cluster1-01
-------------------------------------------------------------------------------
sysconfig: Card in slot 2 (7-1275-0008-46848) is not supported.
sysconfig: slot 12 OK: X2067: Proprietary embedded SAS HBA

cluster1::>
```

```
[[ID081ad1b684830d2ed9a5ab8f5b84d079]]
= system controller config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display System Configuration Information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller config show` command displays system configuration
information for the devices present in the controller. To display more
details, use the `-instance` parameter.

== 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>|local}]`` - Node::
Selects the nodes that match this parameter value.

``[-device <text>]`` - Device::
Selects the configuration information that matches the specified device.

``[-subslot <integer>]`` - Subslot Number::
Selects the configuration information that matches the specified subslot.
```

```
``[-info <text>]`` - Device Info::
Selects the configuration information that matches the specified device
information.

``[-slot <text>]`` - Slot::
Selects the configuration information that matches the specified slot.

== Examples

The following example displays configuration information for slot 1 of the
controller:
```

cluster1::> system controller config show -slot 1

Node: node1
Sub- Device/
Slot slot Information
---- ---- ----------
1 - NVRAM10 HSL
Device Name: Interconnect HBA: Generic OFED Provider
Port Name: ib1a
Default GID: fe80:0000:0000:0000:0000:0000:0000:0104
Base LID: 0x104
Active MTU: 8192
Data Rate: 0 Gb/s (8X)
Link State: DOWN
QSFP Vendor: Amphenol
QSFP Part Number: 112-00436+A0
QSFP Type: Passive Copper 1m ID:00
QSFP Serial Number: APF16130066875
QSFP Vendor: Amphenol
QSFP Part Number: 112-00436+A0
QSFP Type: Passive Copper 1m ID:00
QSFP Serial Number: APF16130066857

cluster1::>

```
[[ID2e817f29d07103eeddadcc928becd04b]]
= system controller config pci show-add-on-devices
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display PCI devices in expansion slots

*Availability:* This command is available to _cluster_ administrators at
```

```
the _admin_ privilege level.

== Description

The `system controller config pci show-add-on-devices` command displays
information about the PCIe devices in I/O expansion slots. The command
displays the following information about the PCIe devices:

* Node
* Model
* Type
* Slot
* Device
* Vendor
* Sub-device ID

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about PCI devices.

``[-node {<nodename>|local}]`` - Node::
Selects the PCI devices that are present in the specified node.

``[-model <text>]`` - Model String::
Selects the PCI devices that are present on the system with the specified
model name.

``[-type <integer>]`` - Device Type::
Selects the PCI devices with the specified device type.

``[-slot-and-sub <text>]`` - PCI Slot Number::
Selects the PCI devices present in the specified slot or slot-subslot
combination.

``[-device <text>]`` - Device::
Selects the PCI devices with the specified device ID.

``[-vendor <text>]`` - Vendor Number::
Selects the PCI devices with the specified vendor ID.

``[-sub-device-id <integer>]`` - Sub Device ID::
```

```
Selects the PCI devices with the specified sub-device ID.

== Examples

The example below displays information about PCI devices found in I/O
expansion slots of all the nodes in the cluster.
```

cluster1::> system controller config pci show-add-on-devices
Node Model Slot Type Device Vendor Sub-Device ID
---------- ---------- ---- ---- ------- ------ ----------
cluster1-01 FAS6240
6 7 0x2532 0x1077 10
5 1 0x1527 0x8086 0
2 7 0x6732 0x15B3 0
3 1 0x8030 0x1077 0
1 2 0x8001 0x11F8 0
15 1 0x10FB 0x8086 0
13 1 0x150E 0x8086 1
7 1 0x1528 0x8086 0
cluster1-02 FAS6240
6 7 0x2532 0x1077 10
5 1 0x1527 0x8086 0
2 7 0x6732 0x15B3 0
3 1 0x8030 0x1077 0
1 2 0x8001 0x11F8 0
15 1 0x10FB 0x8086 0
13 1 0x150E 0x8086 1
7 1 0x1528 0x8086 0
16 entries were displayed.

cluster1::>

```
[[ID528bc95b3eabc143c87c8297ff4471bf]]
= system controller config pci show-hierarchy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display PCI hierarchy

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller config pci show-hierarchy` command displays the PCI
Hierarchy of all PCI devices found in a controller. The command displays
```

```
the following information about the PCI devices:

* Node
* Level
* Device
* Link Capability
* Link Status

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for PCI devices.

``[-node {<nodename>|local}]`` - Node::
Displays the PCI hierarchy of the specified node.

``[-level <integer>]`` - PCI Device Level::
Displays the PCI devices that match the specified level within the PCI
hierarchy.

``[-pci-device <text>]`` - PCI Device::
Displays the PCI devices that match the specified device description.

``[-link-cap <text>]`` - Link Capability::
Displays the PCI devices that match the specified link capability.

``[-link-status <text>]`` - Link Status::
Displays the PCI devices that match the specified link status.

== Examples

The example below displays the PCI hierarchy for all of the nodes in the
cluster.
```

cluster1::> system controller config pci show-hierarchy
PCI Hierarchy

Node: cluster1-01

Level Device Link
----- --------- ------------------------
1 Br[3721](0,3,0): PCI Device 8086:3721 on Controller

```
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
LinkStatus(LkSp(2),LkWd(4),DLAct),
2 Dv[8001](1,0,0): PMC SAS adapter on Controller
LinkCap(MaxLkSp(2),MaxLkWd(8),ASPM(3),L0(3),L1(6),Port(68))
LinkStatus(LkSp(2),LkWd(4),SClk),
1 Br[3722](0,4,0): PCI Device 8086:3722 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),DLAct),
2 Dv[6274](2,0,0): PCI Device 15b3:6274 on Controller
LinkCap(MaxLkSp(1),MaxLkWd(8),ASPM(1),L0(7),L1(7),Port(68))
LinkStatus(LkSp(1),LkWd(4)),
1 Br[3723](0,5,0): PCI Device 8086:3723 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(0)),
1 Br[3b42](0,28,0): PCI Device 8086:3b42 on Controller
LinkCap(MaxLkSp(1),MaxLkWd(4),ASPM(3),L0(4),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk,DLAct),
2 Dv[150e](4,0,0): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
2 Dv[150e](4,0,1): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
2 Dv[150e](4,0,2): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
2 Dv[150e](4,0,3): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
1 Br[3b4a](0,28,4): PCI Device 8086:3b4a on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(4),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk,DLAct),
2 Dv[10d3](5,0,0): Intel 1G NIC on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(1),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk),
1 Br[3b4e](0,28,6): PCI Device 8086:3b4e on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(4),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk,DLAct),
2 Dv[10d3](7,0,0): Intel 1G NIC on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(1),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk),

Node: cluster1-02

Level Device Link
----- --------- -------------------------
1 Br[3721](0,3,0): PCI Device 8086:3721 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
LinkStatus(LkSp(2),LkWd(4),DLAct),
2 Dv[8001](1,0,0): PMC SAS adapter on Controller
LinkCap(MaxLkSp(2),MaxLkWd(8),ASPM(3),L0(3),L1(6),Port(68))
LinkStatus(LkSp(2),LkWd(4),SClk),
1 Br[3722](0,4,0): PCI Device 8086:3722 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
```

LinkStatus(LkSp(1),LkWd(4),DLAct),
2 Dv[6274](2,0,0): PCI Device 15b3:6274 on Controller
LinkCap(MaxLkSp(1),MaxLkWd(8),ASPM(1),L0(7),L1(7),Port(68))
LinkStatus(LkSp(1),LkWd(4)),
1 Br[3723](0,5,0): PCI Device 8086:3723 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(0)),
1 Br[3b42](0,28,0): PCI Device 8086:3b42 on Controller
LinkCap(MaxLkSp(1),MaxLkWd(4),ASPM(3),L0(4),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk,DLAct),
2 Dv[150e](4,0,0): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
2 Dv[150e](4,0,1): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
2 Dv[150e](4,0,2): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
2 Dv[150e](4,0,3): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(4),SClk),
1 Br[3b4a](0,28,4): PCI Device 8086:3b4a on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(4),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk,DLAct),
2 Dv[10d3](5,0,0): Intel 1G NIC on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(1),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk),
1 Br[3b4e](0,28,6): PCI Device 8086:3b4e on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(4),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk,DLAct),
2 Dv[10d3](7,0,0): Intel 1G NIC on Controller
LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(1),L1(6),Port(68))
LinkStatus(LkSp(1),LkWd(1),SClk),
28 entries were displayed.
cluster::>

```
[[IDc2aa43f9529ae525cb40b755d41efbc0]]
= system controller coredump-device show-serial-number
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the coredump device serial number

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `system controller coredump-device show-serial-number` command
displays the serial number of coredump devices. This command is only
applicable to AFF A700 systems. By default, the command displays the
following information about the coredump device:

* Node name
* Display name
* Serial number
* Size
* Device state
* Status

To display more details, use the `-instance` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for the coredump device.

``[-node {<nodename>|local}]`` - Node::
Selects the coredump device that is present on the specified node.

``[-serial-num <text>]`` - Serial Number::
Selects the coredump device with the specified serial number.

``[-vendor-id <Hex Integer>]`` - Vendor ID::
Selects the coredump device with the specified vendor ID.

``[-device-id <Hex Integer>]`` - Device ID::
Selects the coredump device with the specified device ID.

``[-display-name <text>]`` - Display Name::
Selects the coredump device with the specified display name.

``[-unique-name <text>]`` - Unique Name::
Selects the coredump device with the specified unique name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the coredump device with the specified health monitor.

``[-health-status {present|not-present}]`` - Health Status::
```

```
Selects the coredump device with the specified health status.

``[-device-state {good|warn|bad}]`` - Device State::
Selects the coredump device with the specified device state.

``[-size <integer>]`` - Max Memory Size (GB)::
Selects the coredump device with the specified memory size.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the coredump device with the specified health monitor status.

``[-shelf-status {present|not-present}]`` - Shelf Connected Status::
Selects the coredump device with the specified shelf connected status.

== Examples

The following example displays the information of the coredump device
present in all the nodes in a cluster:
```

cluster1::> system controller coredump-device show-serial-number
Size Device
Node FRU/Display Name Serial Number (GB) State Status
---------- ---------- ---------- ------- --------- ------
cluster1-01
Coredump/X9170A A22P5061550000644 1024 good ok
cluster1-02
Coredump/X9170A A22P5061550000137 1024 good ok
2 entries were displayed.

```
The following example displays the detailed information about the coredump
device present in a node:
```

cluster1::> system controller coredump-device show-serial-number -instance -node cluster1-01
Node: cluster1-01
Serial Number: A22P5061550000644
Vendor ID: 1b85
Device ID: 4018
Display Name: Coredump/X9170A
Unique Name: 0X91701190020741OCZ000Z63000001T0003500
Health Monitor Name: controller
Health Status: present
Device State: good
Max Memory Size (GB): 1024
Status: ok
Shelf Connected Status: present

```
```

```
[[ID3edefc1f7804b201a5e15c22a992c2ce]]
= system controller coredump-device show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the coredump device health status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller coredump-device show` command displays details of
the coredump devices present in all the nodes in a cluster. This command
is only applicable to AFF A700 systems. By default, the command displays
the following information about the coredump device:

* Node name
* Display name
* Vendor ID
* Device ID
* Memory size
* Device state
* Health monitor status

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for the coredump device.

``[-node {<nodename>|local}]`` - Node::
Selects the coredump device that is present on the specified node.

``[-serial-num <text>]`` - Serial Number::
Selects the coredump device with the specified serial number.

``[-vendor-id <Hex Integer>]`` - Vendor ID::
Selects the coredump device with the specified vendor ID.
```

```
``[-device-id <Hex Integer>]`` - Device ID::
Selects the coredump device with the specified device ID.

``[-display-name <text>]`` - Display Name::
Selects the coredump device with the specified display name.

``[-unique-name <text>]`` - Unique Name::
Selects the coredump device with the specified unique name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the coredump device with the specified health monitor.

``[-health-status {present|not-present}]`` - Health Status::
Selects the coredump device with the specified health status.

``[-device-state {good|warn|bad}]`` - Device State::
Selects the coredump device with the specified device state.

``[-size <integer>]`` - Max Memory Size (GB)::
Selects the coredump device with the specified memory size.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the coredump device with the specified health monitor status.

``[-shelf-status {present|not-present}]`` - Shelf Connected Status::
Selects the coredump device with the specific shelf connected status.

== Examples

The following example displays the information of the coredump devices
present in all the nodes in a cluster:
```

cluster1::> system controller coredump-device show
Size Device
Node FRU/Display Name Vendor ID Device ID (GB) State Status
---------- --------- --------- --------- ------- --------- ------
cluster1-01
Coredump/X9170A 1b85 4018 1024 good ok
cluster1-02
Coredump/X9170A 1b85 4018 1024 good ok
2 entries were displayed.

```
The example below displays the detailed information about the coredump
device present in a node.
```

cluster1::> system controller coredump-device show -instance -node cluster1-01
Node: cluster1-01
Serial Number: A22P5061550000644
Vendor ID: 1b85
Device ID: 4018
Display Name: Coredump/X9170A
Unique Name: 0X91701190020741OCZ000Z63000001T0003500
Health Monitor Name: controller
Health Status: present
Device State: good
Max Memory Size (GB): 1024
Status: ok
Shelf Connected Status: present

```
[[ID56dad1df9503013f6ff9b0c65ab7f1a1]]
= system controller environment show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the FRUs in the controller

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller environment show` displays information about all
environment FRUs in the cluster. These commands are available for 80xx,
25xx and later systems. Earlier models are not supported. By default, the
command displays the following information about the environment FRUs in
the cluster:

* Node
* FRU name
* FRU state

To display more details, use the -instance parameter.

== Parameters
```

```
{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about the environment FRUs.

``[-node {<nodename>|local}]`` - Node::
Selects information about all the environment FRUs that the specified node
owns.

``[-serial-number <text>]`` - FRU Serial Number::
Selects information about all the environment FRUs with the specified
serial number.

``[-fru-name <text>]`` - FRU Name::
Selects information about the environment FRU with the specified FRU name.

``[-type {controller|psu|fan|dimm|bootmedia|ioxm|nvram|nvdimm}]`` - FRU
Type::
Selects information about all the environment FRUs with the specified FRU
type.

``[-name <text>]`` - Name::
Selects information about all the environment FRUs with the specified
unique name.

``[-state <text>]`` - FRU State::
Selects information about all the environment FRUs with the specified FRU
state.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects information about all the environment FRUs with the specified
health monitor status.

``[-display-name <text>]`` - Display Name for the FRU::
Selects information about all the environment FRUs with the specified
display name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Monitor Name::
Selects information about all the environment FRUs with the specified
monitor.

``[-model <text>]`` - Model Type::
Selects information about all the environment FRUs with the specified FRU
```

```
model.

``[-shared {shared|not_shared}]`` - Shared Resource::
Selects information about all the environment FRUs with the specified
sharing type.

``[-chassis-id <text>]`` - Chassis ID::
Selects information about all the environment FRUs in the specified
chassis.

``[-additional-info <text>]`` - Additional Information About the FRU::
Selects information about all the environment FRU with specified
additional information.

``[-seq-state-cnt <integer>]`` - Count of Same State::
Selects information about all the environment FRU with specified
sequential state count.

== Examples

The following example displays information about all major environment
FRUs in the cluster:
```

cluster1::> system controller environment show

Node FRU Name State
--- --- ---
node1 PSU1 FRU GOOD
node1 PSU2 FRU GOOD
node2 PSU1 FRU GOOD
node2 PSU2 FRU GOOD

```
The following example displays detailed information about a specific
environment FRU:
```

cluster1::> system controller environment show -node node1 -fru-name "PSU1 FRU" -instance
Node: node1
FRU Serial Number: XXT122737891
FRU Name: PSU1 FRU
FRU Type: psu
Name: XXT122737891
FRU State: GOOD
Status: ok
Display Name for the FRU: PSU1 FRU
Monitor Name: controller
Model Type: none
Shared Resource: shared
Chassis ID: 4591227214

Additional Information About the FRU: Part Number: 114-00065+A0
Revision: 020F
Manufacturer: NetApp
FRU Name: PSU

```
[[ID5224f9b9f51334a523d6987c76b4093b]]
= system controller flash-cache show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Flash Cache device status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller flash-cache show` command displays the current
Flash Cache device information.

== 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>|local}]`` - Node::
If this parameter is specified, only status information for the matching
node is displayed.

``[-device-id <integer>]`` - Device ID::
If this parameter is specified, only status information for the matching
device-id is displayed.

``[-slot <integer>]`` - Slot::
If this parameter is specified, only status information for the matching
slot is displayed. Slot can have a format of slot, slot-subslot, or
domain:bus:slot.
```

```
``[-subslot <integer>]`` - Subslot::
If this parameter is specified, only status information for the matching
subslot is displayed.

``[-domain <integer>]`` - Domain::
If this parameter is specified, only status information for the matching
PCI domain is displayed.

``[-bus <integer>]`` - Bus::
If this parameter is specified, only status information for the matching
PCI bus is displayed.

``[-slot-string <text>]`` - Slot String::
If this parameter is specified, only status information for the matching
slot is displayed. Format can be slot or slot-subslot.

``[-device-state
{ok|erasing|erased|failed|removed|online|offline_failed|degraded|offline_t
hreshold}]`` - Device State::
If this parameter is specified, only status information for the matching
device-state is displayed.

``[-model-number <text>]`` - Model Number::
If this parameter is specified, only status information for the matching
model-number is displayed.

``[-part-number <text>]`` - Part Number::
If this parameter is specified, only status information for the matching
part-number is displayed.

``[-serial-number <text>]`` - Serial Number::
If this parameter is specified, only status information for the matching
serial-number is displayed.

``[-firmware-version <text>]`` - Firmware Version::
If this parameter is specified, only status information for the matching
firmware-version is displayed.

``[-firmware-file <text>]`` - Firmware Filename::
If this parameter is specified, firmware-file is the base name of the
firmware file located in disk_fw to update the device firmware.

``[-hardware-revision <text>]`` - Hardware Revision::
If this parameter is specified, only status information for the matching
hardware-revision is displayed.

``[-capacity <integer>]`` - Capacity::
```

```
If this parameter is specified, only status information for the matching
capacity is displayed.

``[-last-change-time <integer>]`` - Time Last State Change::
If this parameter is specified, only status information for the matching
last-change-time is displayed.

``[-service-time <integer>]`` - Service Time::
If this parameter is specified, only status information for the matching
service-time is displayed.

``[-percent-online <integer>]`` - Percent Online::
If this parameter is specified, only status information for the matching
percent-online is displayed.

``[-average-erase-cycle-count <integer>]`` - Avg Erase Cycle Count::
If this parameter is specified, only status information for the matching
average-erase-cycle-count is displayed.

``[-threshold-profile <text>]`` - Threshold Profile::
If this parameter is specified, only status information for the matching
threshold-profile is displayed.

== Examples

The following example displays the current state of all Flash Cache
devices:
```

```
cluster1::> system controller flash-cache show
Device Model Part Serial Firmware Capacity Device
Node ID Slot Number Number Number Version (GB) State
---------- ------ ---- ------ --------- -------------------- -------- -------- ------
node1
0 6-1 X9172A 119-00209 A22P7061550000004 NA00 4096 ok
1 6-2 X9170A 119-00207 A22P5061550000135 NA00 1024 ok
node2
0 6-1 X9172A 119-00209 A22P7061550000007 NA00 4096 ok
1 6-2 X9170A 119-00207 A22P5061550000091 NA00 1024 ok
4 entries were displayed.
```

```
[[ID9fc4eb815e3e2c683f96e66cd34bdbd1]]
= system controller flash-cache secure-erase run
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Perform a secure-erase operation on the targeted devices

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller flash-cache secure-erase run` command securely
erases the given Flash Cache device.

== Parameters

``-node {<nodename>|local}`` - Node::
Selects the node of the specified Flash Cache devices.

``-device-id <integer>`` - Device Id::
Selects the device-id of the specified Flash Cache devices.

== Examples

The following example securely erases the selected Flash Cache device:
```

cluster1::> system controller flash-cache secure-erase -node node1 -device-id 0

```
[[ID8f3b42d174d21eb2e1573f463fba2621]]
= system controller flash-cache secure-erase show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Flash Cache card status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller flash-cache secure-erase show` command displays the
current Flash Cache device secure-erase status.

== 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>|local}]`` - Node::
If this parameter is specified, only status information for the matching
node is displayed.

``[-device-id <integer>]`` - Device Id::
If this parameter is specified, only status information for the matching
device-id is displayed.

``[-device-state {ok|erasing|erased|failed|removed}]`` - Device State::
If this parameter is specified, only status information for the matching
device-state is displayed.

== Examples

The following example displays the current state of all the Flash Cache
devices:
```

cluster1::> system controller flash-cache secure-erase show
Node Device ID Slot Device State
--------- --------- ---- ------------
node1
0 6-1 ok
1 6-2 erasing
node2
0 6-1 erased
1 6-2 ok
4 entries were displayed.

```
[[ID52c419ece565900c9c957de634676b59]]
= system controller fru show-manufacturing-info
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display manufacturing information of FRUs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller fru show-manufacturing-info` command displays
manufacturing information for field replaceable units (FRUs) installed in
the system. The information includes FRU-description, serial number, part
number, and revision number. To display more details, use the `-instance`
parameter.

== 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  ]`` }::
Displays detailed information about the installed FRUs in the system.

``[-node {<nodename>|local}]`` - Node::
Selects a specific node's installed FRUs.

``[-system-sn <text>]`` - System Serial Number::
Selects information about installed FRUs with the specified system serial
```

```
number.

``[-model-name <text>]`` - Model Name::
Selects information about installed FRUs with the specified model name.

``[-system-id <text>]`` - System ID::
Selects information about installed FRUs with the specified system ID.

``[-kernel-version <text>]`` - Kernel Version::
Selects information about installed FRUs with the specified kernel
version.

``[-firmware-release <text>]`` - Firmware Release::
Selects information about installed FRUs with the specified firmware
release.

``[-description <text>]`` - FRU Description::
Selects information about installed FRUs with the specified FRU
description.

``[-vendor-id <text>]`` - FRU Vendor ID::
Selects information about the FRU with the vendor identification.

``[-serial-number <text>]`` - FRU Serial Number::
Selects information about the FRU with the specified serial number.

``[-part-number <text>]`` - FRU Part Number::
Selects information about the FRU with the specified part number.

``[-revision <text>]`` - FRU Revision of Part Number::
Selects information about the FRU with the specified revision.

``[-manufacturer <text>]`` - FRU Manufacturer::
Selects information about the FRU with the specified manufacturer.

``[-manufacture-date <text>]`` - FRU Manufacturing Date::
Selects information about the FRU with the specified manufacture date.

``[-product-id <text>]`` - FRU Product Identifier::
Selects information about the FRU with the specified product ID.

``[-firmware-version <text>]`` - FRU Firmware Version::
Selects information about the FRU with the specified firmware version.

== Examples

The following example displays all installed FRUs in the system:
```

```
cluster1::> system controller fru show-manufacturing-info
Node: platsw-lodi-1-01
System Serial Number: 791541000047
Model Name: FAS9040
System ID: 0537024373
Firmware release: 10.0X18
Kernel Version: NetApp Release sysMman_3887886_1608151712: Mon Aug 15
15:54:00 PDT 2016
FRU Description FRU Serial Number FRU Part Number FRU Rev.
----------------------- ----------------------- ------------------ ----------
Mother Board 031537000390 111-02419 40
Chassis 031536000252 111-02392 40
DIMM-1 CE-01-1510-02A8DC73 SHB722G4LML23P2-SB -
DIMM-3 CE-01-1510-02A8DCCC SHB722G4LML23P2-SB -
DIMM-8 CE-01-1510-02A8DE54 SHB722G4LML23P2-SB -
DIMM-9 CE-01-1510-02A8DE1C SHB722G4LML23P2-SB -
DIMM-11 CE-01-1510-02A8DF42 SHB722G4LML23P2-SB -
DIMM-16 CE-01-1510-02A8DD9B SHB722G4LML23P2-SB -
FAN1 031534001263 441-00058 40
FAN2 031534001292 441-00058 40
FAN3 031534001213 441-00058 40
PSU1 PSD092153200591 114-00146 40
PSU3 PSD092153200700 114-00146 40
mSATA boot0 1439100B02C3 - MU03
1/10 Gigabit Ethernet Controller IX4-T 031538000121 111-02399 40
QLogic 8324 10-Gigabit Ethernet Controller 031535000664 111-02397 40
NVRAM10 031537000846 111-02394 40
NVRAM10 BATT 31534000932 NetApp, Inc. 111-02591
NVRAM10 DIMM CE-01-1510-02A8DC03 SHB722G4LML23P2-SB -
PMC-Sierra PM8072 (111-02396) 031537000246 111-02396 41
PMC-Sierra PM8072 (111-02396) 031537000246 111-02396 41
PMC-Sierra PM8072 (111-02396) 031537000246 111-02396 41
PMC-Sierra PM8072 (111-02396) 031537000246 111-02396 41
PMC-Sierra PM8072 (111-02396) 031537000179 111-02396 41
Disk Serial Number PNHH1J0B X421_HCOBD450A10 -
Disk Serial Number PNHH2BKB X421_HCOBD450A10 -
Disk Serial Number PNHJPZ8B X421_HCOBD450A10 -
Disk Serial Number PNHG6SKB X421_HCOBD450A10 -
Disk Serial Number PNHKJYTB X421_HCOBD450A10 -
Disk Serial Number PNHSVMEY X421_HCOBD450A10 -
Disk Serial Number PNHT8KWY X421_HCOBD450A10 -
Disk Serial Number PNHSVL0Y X421_HCOBD450A10 -
Disk Serial Number PNHG5RHB X421_HCOBD450A10 -
Disk Serial Number PNHEXLWB X421_HCOBD450A10 -
Disk Serial Number PNHT8LJY X421_HCOBD450A10 -
Disk Serial Number PNHSVKZY X421_HCOBD450A10 -
PMC-Sierra PM8072 (111-02396) 031537000179 111-02396 41
PMC-Sierra PM8072 (111-02396) 031537000179 111-02396 41
PMC-Sierra PM8072 (111-02396) 031537000179 111-02396 41
DS2246 6000113106 0190 -
DS2246-Pwr-Supply XXT111825308 114-00065+A0 9C
DS2246-Pwr-Supply XXT111825314 114-00065+A0 9C
DS2246-MODULE 8000675532 111-00690+A3 23
DS2246-MODULE 8000751790 111-00690+A3 23
```

DS2246-CABLE 512130075 112-00430+A0 -
DS2246-CABLE - - -
DS2246-CABLE 512130118 112-00430+A0 -
DS2246-CABLE - - -
49 entries were displayed.

```
[[ID12b0b00e350ffbbab5039c77b0795320]]
= system controller fru show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Information About the FRUs in the Controller

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller fru show` command displays information about all
the controller specific Field Replaceable Units (FRUs) in the cluster.
These commands are available for 80xx, 25xx and later systems. Earlier
models are not supported. By default, the command displays the following
information about all the FRUs in the cluster:

* Node
* FRU name
* Health monitor subsystem
* Health monitor status

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about the controller specific FRUs in the
cluster.

``[-node {<nodename>|local}]`` - Node::
Selects information about the FRUs in the specified node.

``[-subsystem <Subsystem>]`` - Subsystem::
```

```
Selects information about the FRUs of the specified subsystem.

``[-serial-number <text>]`` - FRU Serial Number::
Selects information about the FRU with the specified serial number.

``[-fru-name <text>]`` - Name of the FRU::
Selects information about the FRU with the specified FRU name.

``[-type {controller|psu|fan|dimm|bootmedia|ioxm|nvram|nvdimm}]`` - FRU
Type::
Selects information about the FRU with the specified FRU type.

``[-name <text>]`` - FRU Name::
Selects information about the FRU with the specified unique name.

``[-state <text>]`` - FRU State::
Selects information about the FRU with the specified state.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects information about the FRU with the specified health monitor
status.

``[-display-name <text>]`` - Display Name for the Fru::
Selects information about the FRU with the specified display name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Monitor Name::
Selects information about the FRU with the specified health monitor type.

``[-model <text>]`` - Model Type::
Selects information about the FRU with the specified model.

``[-chassis-id <text>]`` - Chassis ID::
Selects information about the FRU with the specified chassis ID.

``[-location <text>]`` - Location of the FRU::
Selects information about the FRU with the specified FRU location.

``[-additional-info <text>]`` - Additional Information About the FRU::
Selects information about the FRU with the specified additional
information.

== Examples

The example below displays information about all controller specific FRUs
```

```
in the cluster.
```

cluster1::> system controller fru show
Node FRU Name Subsystem Status
--- --- --- ---
node1 PSU1 FRU Environment ok
node1 PSU2 FRU Environment ok
node1 DIMM-NV1 Memory ok
node1 DIMM-1 Memory ok
node1 Micron Technology 0x655 (ad.0) Motherboard ok
node2 PSU1 FRU Environment ok
node2 PSU2 FRU Environment ok
node2 DIMM-NV1 Memory ok
node2 DIMM-1 Memory ok
node2 Micron Technology 0x655 (ad.0) Motherboard ok
10 entries were displayed.

```
The example below displays information about the specific FRU.
```

cluster1::> system controller fru show -instance -serial-number AD-01-1306-2EA01E9A
Node: node1
Subsystem: Memory
FRU Serial Number: AD-01-1306-2EA01E9A
Name of the FRU: DIMM-1
FRU Type: dimm
FRU Name: DIMM-1
FRU State: ok
Status: ok
Display Name for the Fru: DIMM-1
Monitor Name: controller
Model Type: none
Chassis ID: 4591227214
Location of the FRU: Memory Slot: 1
Additional Information About the FRU: Part No: HMT82GV7MMR4A-H9

```
[[ID426af8597670ff6890f48d412386b6e1]]
= system controller fru led disable-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Turn off all the LEDs Data Ontap has lit

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.
```

```
== Description

The `system controller fru led disable-all` command turns off all the
controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed
to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which
there are corresponding fault LEDs. In addition, there is a summary FRU
fault LED on the external face-plate of both the controller and IOXM;
labeled with a "!". A summary fault LED will be on when any of the
internal FRU fault LEDs are on. Only the controller and IOXM internal FRU
fault LEDs can be controlled by the end-user. The summary fault LEDs are
turned on and off based on the simple policy described above. If you want
to turn off the summary fault LED, you must turn off all internal FRU
fault LEDs.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the
system are FRU fault LEDs. Externally visible fault LEDs are labeled with
a "!" and internal FRU fault LEDs remain on, even when the controller or
IOXM is removed from the chassis. In addition, internal FRU fault LEDs
will remain on until explicitly turned off by the end-user, even after a
FRU has been replaced.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs,
cards in PCI slots, boot media devices, NV batteries and coin cell
batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is
the number of FRUs of that particular type that exist in the controller or
IOXM. Both controller and IOXM have a FRU map label for use in physically
locating internal FRUs. The FRU ID/slot tuple used by the
xref:{relative_path}system-controller-fru-led-show.html[system controller
fru led show] command matches that specified on the FRU map label.

== Examples

Turn off all FRU fault LEDs.
```

cluster1::*> system controller fru led disable-all
14 entries were modified.

```
== Related Links

* xref:{relative_path}system-controller-fru-led-show.html[system
controller fru led show]
```

```
[[ID64c2bd1451f0f6e35a461ede5a308073]]
= system controller fru led enable-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Light all the LEDs

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller fru led enable-all` command turns on all the
controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed
to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which
there are corresponding fault LEDs. In addition, there is a summary FRU
fault LED on the external face-plate of both the controller and IOXM;
labeled with a "!". A summary fault LED will be on when any of the
internal FRU fault LEDs are on. Only the controller and IOXM internal FRU
fault LEDs can be controlled by the end-user. The summary fault LEDs are
turned on and off based on the simple policy described above. If you want
to turn off the summary fault LED, you must turn off all internal FRU
fault LEDs.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the
system are FRU fault LEDs. Externally visible fault LEDs are labeled with
a "!" and internal FRU fault LEDs remain on, even when the controller or
IOXM is removed from the chassis. In addition, internal FRU fault LEDs
will remain on until explicitly turned off by the end-user, even after a
FRU has been replaced.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs,
cards in PCI slots, boot media devices, NV batteries and coin cell
batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is
the number of FRUs of that particular type that exist in the controller or
IOXM. Both controller and IOXM have a FRU map label for use in physically
locating internal FRUs. The FRU ID/slot tuple used by the
```

```
xref:{relative_path}system-controller-fru-led-show.html[system controller
fru led show] command matches that specified on the FRU map label.

== Examples

Turn on all FRU fault LEDs.
```

cluster1::*> system controller fru led enable-all
14 entries were modified.

```
== Related Links

* xref:{relative_path}system-controller-fru-led-show.html[system
controller fru led show]




[[ID58a828422aa6c7b178551bbf3a8d3e40]]
= system controller fru led modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the status of FRU LEDs

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller fru led modify` command modifies the current state
of the controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed
to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which
there are corresponding fault LEDs. In addition, there is a summary FRU
fault LED on the external face-plate of both the controller and IOXM;
labeled with a "!". A summary fault LED will be on when any of the
internal FRU fault LEDs are on. Only the controller and IOXM internal FRU
fault LEDs can be controlled by the end-user. The summary fault LEDs are
turned on and off based on the simple policy described above. If you want
```

```
to turn off the summary fault LED, you must turn off all internal FRU
fault LEDs.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the
system are FRU fault LEDs. Externally visible fault LEDs are labeled with
a "!" and internal FRU fault LEDs remain on, even when the controller or
IOXM is removed from the chassis. In addition, internal FRU fault LEDs
will remain on until explicitly turned off by the end-user, even after a
FRU has been replaced.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs,
cards in PCI slots, boot media devices, NV batteries and coin cell
batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is
the number of FRUs of that particular type that exist in the controller or
IOXM. Both controller and IOXM have a FRU map label for use in physically
locating internal FRUs. The FRU ID/slot tuple used by the
xref:{relative_path}system-controller-fru-led-show.html[system controller
fru led show] command matches that specified on the FRU map label.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects FRU fault LEDs on the specified nodes.

``-fru-id <FRU LED key>`` - FRU ID [.advaced-priv]#(privilege:
advanced)#::
Selects the FRU fault LEDs that match the specified FRU type.

``-fru-slot <integer>`` - FRU Slot [.advaced-priv]#(privilege:
advanced)#::
Selects the FRU fault LEDs that match the specified slot.

``[-fru-state {on|off|unknown}]`` - FRU State [.advaced-priv]#(privilege:
advanced)#::
Specifies the target state for the FRU fault LED.

== Examples

Turn off DIMM 3's FRU fault LED.
```

cluster1::*> system controller fru led modify -node node1 -fru-id dimm -fru-slot 3 -fru-state off

```
The example below turns on all PCI FRU fault LEDs.
```

cluster1::*> system controller fru led modify -node node1 -fru-id pci -fru-slot * -fru-state on

```
== Related Links

* xref:{relative_path}system-controller-fru-led-show.html[system
controller fru led show]




[[ID6754f6ab43d9aa892301e8a00f3e203f]]
= system controller fru led show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the status of FRU LEDs

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller fru led show` command displays information about
the current state of the controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed
to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which
there are corresponding fault LEDs. In addition, there is a summary FRU
fault LED on the external face-plate of both the controller and IOXM;
labeled with a "!". A summary fault LED will be on when any of the
internal FRU fault LEDs are on.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the
system are FRU fault LEDs. Externally visible fault LEDs are labeled with
a "!" and internal FRU fault LEDs remain on, even when the controller or
IOXM is removed from the chassis.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs,
cards in PCI slots, boot media devices, NV batteries and coin cell
batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is
the number of FRUs of that particular type that exist in the controller or
```

```
IOXM. Both controller and IOXM have a FRU map label for use in physically
locating internal FRUs. The FRU ID/slot tuple used by the `system
controller fru led show` command matches that specified on the FRU map
label.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects FRU fault LEDs on the specified nodes.

``[-fru-id <FRU LED key>]`` - FRU ID [.advaced-priv]#(privilege:
advanced)#::
Selects the FRU fault LEDs that match the specified FRU type.

``[-fru-slot <integer>]`` - FRU Slot [.advaced-priv]#(privilege:
advanced)#::
Selects the FRU fault LEDs that match the specified slot.

``[-fru-bay <text>]`` - FRU Bay [.advaced-priv]#(privilege: advanced)#::
Selects the FRU fault LEDs that match the specified bay.

``[-fru-state {on|off|unknown}]`` - FRU State [.advaced-priv]#(privilege:
advanced)#::
Selects the FRU fault LEDs that match the specified status.

``[-lit-by <text>]`` - Lit By [.advaced-priv]#(privilege: advanced)#::
Selects the FRU fault LEDs that were lit by the specified source.

== Examples

List the current state of all FRU fault LEDs.
```

cluster1::*> system controller fru led show
Node FRU Type Bay Slot State Lit By
--- --- --- --- --- ---
host1
controller A 1 on SP

ioxm B 1 off -
pci - 1 off -
pci - 2 off -
pci - 3 off -
pci - 4 off -
pci - 5 off -
pci - 6 off -
dimm-nv - 1 off -
dimm-nv - 2 off -
dimm - 1 off -
dimm - 2 off -
dimm - 3 off -
dimm - 4 off -
identify - 1 off -
5 entries were displayed.

```
The example below displays the status of only a specific FRU.
```

cluster1::*> system controller fru led show -node host1 -fru-id controller -fru-slot 1
Node FRU Type Bay Slot State Lit By
--- --- --- --- --- ---
host1
controller A 1 off -

```
[[ID7b7e557dc59cadd07f4a21ad3c36d6a1]]
= system controller ioxm show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays IOXM Device Health Status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller ioxm show` command displays the details of the IO
expansion modules (IOXMs) that are connected to the nodes in a cluster.
These commands are available for 80xx, 25xx and later systems. Earlier
models are not supported. By default, the command displays the following
information about the IOXMs:

* Node name
* Display name
* Is IOXM present?
```

```
* Power status
* Health monitor status

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for all the IOXMs.

``[-node {<nodename>|local}]`` - Node::
Selects the IOXM that is connected to the specified node.

``[-chassis-config {c-i|c-c|c-b}]`` - Controller-IOXM or Controller-
Controller or Controller-Blank::
Selects the IOXMs with the specified chassis configuration.

``[-is-present {present|not-present}]`` - IOXM Presence::
Selects the IOXMs that are connected and detected (`present` ) or
connected but not detected (`not-present` ).

``[-power {good|bad}]`` - Power to IOXM::
Selects the IOXMs with the specified power state.

``[-display-name <text>]`` - Display Name::
Selects the IOXMs with the specified display name.

``[-unique-name <text>]`` - Unique Name::
Selects the IOXM with the specified unique name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the IOXMs with the specified health monitor.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` - IOXM
Health::
Selects the IOXMs with the specified health monitor status.

== Examples

The example below displays the information of all the IOXMs that are
connected to the nodes in a cluster.
```

```
cluster1::> system controller ioxm show

Node Display Name Is-Present? Power Status
------------- ------------ ----------- --------- ------
node1 IOXM present good ok
node2 IOXM present good ok
```

```
The example below displays detailed information of an IOXM that is
connected to a node.
```

```
cluster1::> system controller ioxm show -instance -node node1
Node: node1
Controller-IOXM or Controller-Controller or Controller-Blank: c-i
IOXM Presence: present
Power to IOXM: good
Display Name: node1/IOXM
Unique Name: 8006459930
Health Monitor Name: controller
IOXM Health: ok
```

```
[[IDd1f4aa42fbe09c962962858d3b5e3989]]
= system controller location-led modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the location LED state of a controller

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller location-led modify` command modifies the current
state of the location LED. When lit, the location LED can help you find
the controller in the data center.

There is a blue location LED on every controller and on the front of the
chassis. When you turn on the location LED for either controller, the
chassis location LED automatically turns on. When both controller location
LEDs are off, the chassis location LED automatically turns off.

After the location LED is turned on, it stays illuminated for 30 minutes
and then automatically shuts off.

== Parameters

``-node {<nodename>|local}`` - Node::
Selects the location LED on the specified filers.

``[-state {on|off}]`` - LED State::
Modifies the state of the location LED on the filer.

== Examples

The following example turns on the location LED:
```

cluster1::*> system controller location-led modify -node node1 -state on

```
Turn off Location LED.
```

cluster1::*> system controller location-led modify -node node1 -state off

```
[[ID566591d2e803f0f374d9991e650faa74]]
= system controller location-led show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the location LED state on controllers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller location-led show` command shows the current state
of the location LED. When lit, the location LED can help you find the
controller in the data center.

There is a blue location LED on every controller and on the front of the
chassis. When you turn on the location LED for either controller, the
chassis location LED automatically turns on. When both controller location
LEDs are off, the chassis location LED automatically turns off.

After the location LED is turned on, it stays illuminated for 30 minutes
and then automatically shuts off.

== 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>|local}]`` - Node::
Selects the location LED on the specified filers.

``[-state {on|off}]`` - LED State::
Displays the location LED's status.

== Examples

The following example lists the current state of the location LED:
```

```
cluster1::*> system controller location-led show
Node Location LED State
-------------- -------------------
node1 Off
node2 Off
```

```
[[IDfa233560ec36dc21010ad6aef0b09021]]
= system controller memory dimm show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Memory DIMM Table

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller memory dimm show` command displays information
about the DIMMs in all the nodes in the cluster. These commands are
available for 80xx, 25xx and later systems. Earlier models are not
supported. By default, the command displays the following information
about all the DIMMs in the cluster:

* Node
* DIMM name
* Uncorrectable ECC error count
* Correctable ECC error count
* CECC Alert Method
* CPU socket
* Channel
* Slot number
* Health monitor status
* Failure reason

To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about the DIMMs in all the controllers in
```

```
the cluster.

``[-node {<nodename>|local}]`` - Node::
Selects information about the DIMMs in the specified node.

``[-pds-id <integer>]`` - DIMM ID::
Selects information about the DIMMs with the specified DIMM ID.

``[-slotname <text>]`` - Slot Name::
Selects information about the DIMMs with the specified slot name.

``[-socket <integer>]`` - CPU Socket::
Selects information about the DIMMs with the specified socket ID.

``[-channel <integer>]`` - Channel::
Selects information about the DIMMs with the specified channel number.

``[-slot-no <integer>]`` - Slot Number on a Channel::
Selects information about the DIMMs with the specified slot number.

``[-serial <text>]`` - Serial Number::
Selects information about the DIMMs with the specified serial number.

``[-part-no <text>]`` - Part Number::
Selects information about the DIMMs with the specified part number.

``[-cecc-count <integer>]`` - Correctable ECC Error Count::
Selects information about the DIMMs with the specified correctable ECC
error count.

``[-uecc-count <integer>]`` - Uncorrectable ECC Error Count::
Selects information about the DIMMs with the specified uncorrectable ECC
error count.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects information about the DIMMs with the specified health monitor.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects information about the DIMMs with the specified health monitor
status.

``[-name <text>]`` - Unique Name of DIMM::
Selects information about the DIMMs with the specified unique name.
```

```
``[-display-name <text>]`` - Display Name for the DIMM::
Selects information about the DIMMs with the specified display name.

``[-alt-cecc-method <text>]`` - CECC Alert Method::
Selects information about the DIMMs with the specified CECC error alert
method.

``[-alt-cecc-dimm {true|false}]`` - Replace DIMM::
Selects information about the DIMMs with the specified replace DIMM value.

``[-failure-reason <text>]`` - Failure Reason::
Selects information about the DIMMs with the specified failure reason.

== Examples

The example below displays information about the DIMMs in all the nodes in
the cluster.
```

```
cluster1::> system controller memory dimm show
DIMM UECC CECC Alert CPU Slot Failure
Node Name Count Count Method Socket Channel Number Status Reason
-------------- ------- -------- -------- ------ ------ ------- ------ ------- --------
node1 DIMM-1 0 0 bucket 0 0 0 unknown disabled
DIMM-NV1 0 0 bucket 0 1 1 ok
node2 DIMM-1 1 0 bucket 0 0 0 ok
node3 DIMM-NV1 0 0 bucket 0 1 1 ok
4 entries were displayed.
```

```
The example below displays detailed information about a specific DIMM in a
specific controller.
```

```
cluster1::> system controller memory dimm show -instance -node node1 -pds-id 1
Node: node1
DIMM ID: 1
Slot Name: DIMM-1
CPU Socket: 0
Channel: 0
Slot Number on a Channel: 0
Serial Number: AD-01-1306-2EA01E9A
Part Number: HMT82GV7MMR4A-H9
Correctable ECC Error Count: 0
Uncorrectable ECC Error Count: 0
Health Monitor Name: controller
Status: unknown
Unique Name of DIMM: DIMM-1
Display Name for the DIMM: DIMM-1
CECC Alert Method: bucket
Replace DIMM: false
Failure Reason: disabled
```

```
[[ID07836d4fe1ccfc2b55530ec4634c5fb1]]
= system controller nvram-bb-threshold show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the controller NVRAM bad block threshold

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller nvram-bb-threshold show` command displays the
threshold for the NVRAM bad block counts for a node.




[[IDcf21e383e679c9532e54f655f11ab4f6]]
= system controller pci show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the PCI Device Table

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller pci show` command displays details of the PCI
devices present in all of the nodes in a cluster. These commands are
available for 80xx, 25xx and later systems. Earlier models are not
supported. By default, the command displays the following information
about the PCI devices:

* Node name
* Display name
* Correctable error count
* Functional link width
```

```
* Functional link speed
* Health monitor status

 To display more details, use the -instance parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information for all of the PCI devices.

``[-node {<nodename>|local}]`` - Node::
Selects the PCI devices that are present in the specified node.

``[-bus-number <integer>]`` - Bus Number::
Selects the PCI devices with the specified bus number.

``[-device-number <integer>]`` - Device Number::
Selects the PCI devices with the specified device number.

``[-function-number <integer>]`` - Function Number::
Selects the PCI devices with the specified function number.

``[-slot-number <integer>]`` - Slot Info::
Selects the PCI devices with the specified slot number.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the PCI devices monitored by the specified health monitor.

``[-vendor-id <Hex Integer>]`` - Vendor ID::
Selects the PCI devices with the specified vendor ID.

``[-device-id <Hex Integer>]`` - Device ID::
Selects the PCI devices with the specified device ID.

``[-physical-link-width <integer>]`` - Physical Link Width::
Selects the PCI devices with the specified physical link width.

``[-functional-link-width <integer>]`` - Functional Link Width::
Selects the PCI devices with the specified functional link width.

``[-physical-link-speed <text>]`` - Physical Link Speed(GT/s)::
Selects the PCI devices with the specified physical link speed.
```

```
``[-functional-link-speed <text>]`` - Functional Link Speed(GT/s)::
Selects the PCI devices with the specified functional link speed.

``[-unique-name <text>]`` - Unique Name::
Selects the PCI devices with the specified unique name.

``[-corr-err-count <integer>]`` - Correctable Error Count::
Selects the PCI devices with the specified correctable error count.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the PCI devices with the specified health monitor status.

``[-display-name <text>]`` - Display Name::
Selects the PCI devices with the specified display name.

``[-cerr-diff <integer>]`` - Correctable Error Difference::
Selects the PCI devices with the specified difference in correctable error
count.

== Examples

The example below displays the information about the PCIe devices present
in all of the nodes in the cluster.
```

cluster1::> system controller pci show
Display Correctable Functional Functional
Node Name Error Count Link Width Link Speed Status
---------- ---------- ---------- ---------- ---------- ------
cluster1-01 Ontap PCI Device 0 0 4 5GT/s ok
cluster1-02 Ontap PCI Device 4 0 4 5GT/s ok

```
The example below displays detailed information about a PCIe device in a
node.
```

cluster1::> system controller pcie show -instance -node cluster1-01 -bus-number 1
Node: cluster1-01
Bus Number: 1
Device Number: 0
Function Number: 0
Slot Info: 0
Health Monitor Name: controller
Vendor ID: 11f8
Device ID: 8001
Physical Link Width: 4
Functional Link Width: 4

Physical Link Speed(GT/s): 5GT/s
Functional Link Speed(GT/s): 5GT/s
Unique Name: ontap0@pci0:1:0:0
Correctable Error Count: 0
Status: ok
Display Name: Ontap PCI Device 0
Correctable Error Difference: 0

```
[[ID114370234a2fbe68732e6c6d3f466e56]]
= system controller pcicerr threshold modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the Node PCIe error alert threshold

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller pcicerr threshold modify` command modifies node-
wide PCIe correctable error threshold counts in the cluster.

== Parameters

``[-pcie-cerr-threshold <integer>]`` - Corr. Error Limit::
The PCIe error threshold count that would trigger an alert if exceeded.

``[-nvram-bb-threshold <integer>]`` - NVRAM Bad Block limit::
The NVRAM bad block threshold count that would trigger an alert if
exceeded.

== Examples

The example below displays the information about setting node-wide PCIe
error threshold count in the cluster:
```

cluster1::> system controller threshold modify -pcie-cerr-threshold 100

```
[[IDbdbb597e057f975863a0bd75b15d6899]]
= system controller pcicerr threshold show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Node PCIe error alert threshold

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller pcicerr threshold show` command displays
information about node-wide PCIe correctable error threshold counts in the
cluster.

== Examples

The example below displays the information about node-wide PCIe error
threshold count in the cluster:
```

cluster1::> system controller pcicerr threshold show

PCIe Error Threshold

200

```
[[ID67c1819dfb82feb770d67be7498a6a06]]
= system controller platform-capability show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display platform capabilities

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `system controller platform-capability show` command displays
information about all platform capabilities for each controller in the
cluster. By default, the command displays the following information about
all controllers in the cluster:

* Controller Name
* Capability ID
* Capability Supported?
* Capability Name



== 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  ]`` }::
Displays detailed information about all controllers in the cluster.

``[-node {<nodename>|local}]`` - Node::
Selects information about the specified controller.

``[-capability-id <integer>]`` - Capability ID::
Selects the desired capability ID.

``[-supported <text>]`` - Supported?::
Selects the desired capability support state (true or false).

``[-name <text>]`` - Capability Name::
Selects the desired capability name.

== Examples

The following example displays platform capability information for the
controller:
```

cluster1::> system controller platform-capability show
Node Capability ID Supported? Capability Name
--- --- --- ---
or-099-diag-01
0 false CAP_CMCI_ENABLED
1 false CAP_HA_CONFIG_ONLY
2 true CAP_SUPPORT_CARD_FRU
3 true CAP_SCORPIO_EN

4 false CAP_NVD_EN
5 false CAP_ENABLE_HPET
6 false CAP_VERIFY_ACPI_TABLE
7 entries were displayed.

```
[[IDb070e8b1f1167bce6994107d4e69433e]]
= system controller replace cancel
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cancel ongoing controller replacement

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller replace cancel` command is used to cancel a
controller replacement that is in paused state (paused-on-request, paused-
on-error or paused-for-intervention). The update cannot be canceled if it
is not in a paused state.

== Examples

The following example displays a cancel operation:
```

cluster1::> system controller replace cancel
Warning: The controller replacement will be canceled and any changes will have to be reverted manually.
Do you want to continue? {y|n}: y
Controller replacement canceled successfully.

```
[[IDbb335b8da988370b6bec38f80e04640d]]
= system controller replace pause
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Pause ongoing controller replacement

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller replace pause` command is used to pause a currently
running replacement. The operation pauses at the next predefined update
stage (for example, after finishing the current task it will pause the
next restartable task) which might take some time to reach. When the
update reaches the pause point, it transitions into the pause-on-request
state.

== Examples

The following example displays pause operation:
```

cluster1::> system controller replace pause

A pause requested for Controller Replacement operation.
The current task will continue and the next restartable task will be paused.

```
[[IDd86e823dd08d80017f0c753517d20bfb]]
= system controller replace resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume paused controller replacement

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller replace resume` command is used to resume an update
that is currently in one of paused-on-request, paused-on-error or paused-
for-intervention states. If the update is not paused then an error is
returned.

== Examples

The following example shows a resume operation:
```

cluster1::> system controller replace resume

Controller replacement resumed successfully.

```
[[IDbf9ae6e1ef2e5bc69cc707ffb774c5f2]]
= system controller replace show-details
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display detailed status of controller replacement

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller replace show-details` command displays detailed
information about the currently running and previously run non-disruptive
```

```
controller replacement operations. The command displays the following
information:

* Phase
* Node
* Task name
* Task status
* Error message



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-operation-identifier {None|Controller-replacement}]`` - Operation
Identifier [.advaced-priv]#(privilege: advanced)#::
Specifies the NDO operation identifier.

``[-task-identifier <integer>]`` - Task Identifier [.advaced-
priv]#(privilege: advanced)#::
Specifies the identification number of the task.

``[-node <nodename>]`` - Node That Performs Operation [.advaced-
priv]#(privilege: advanced)#::
Specifies the node that is to be replaced.

``[-task-phase {None|Initialization|Precheck|Commit|Collect-
info|Preop|Resource-release|Verification|Resource-
regain|Postop|Uncommit|Postcheck|Cleanup}]`` - Task Phase [.advaced-
priv]#(privilege: advanced)#::
Specifies the phase of the operation.

``[-task-name <text>]`` - Name of the Task [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the task.

``[-task-status {None|Ready-to-start|In-progress|Waiting|Paused-on-
error||Paused-for-intervention|Paused-on-request|Completed-on-first-
node|Completed|Failed|Pause_req|Canceled}]`` - Status of the Task
```

```
[.advaced-priv]#(privilege: advanced)#::
Specifies the status of the task.

``[-task-error <text>]`` - Error During the Task Execution [.advaced-
priv]#(privilege: advanced)#::
Specifies the error occurred.

``[-task-recovery-action <text>]`` - Action to Recover from Error
[.advaced-priv]#(privilege: advanced)#::
Specifies the action to be taken in case of error.

== Examples

The following example displays detailed information about the non-
disruptive replacement operation:
```

cluster1::*> system controller replace show-details
Task Phase Node Task Name Operation-State
--- --- --- ---
Precheck node1 Cluster Health Check Completed
MCC Cluster Check Completed
Aggr Relocation Completed
Status Check
Model Name Check Completed
Cluster Quorum Check Completed
Image Version Check Completed
HA Status Check Completed
Aggregate Status Completed
Check
Disk Status Check Completed
Data LIF Status Check Completed
Cluster LIF Status Completed
Check
ASUP Status Check Completed
CPU Utilization Check Completed
Aggr Reconstruction Completed
Check
Node Affinity Job Completed
Check
Collect-info node1 Verify Details Paused-for-intervention
node2 Verify Details Paused-for-intervention
17 entries were displayed.

```
[[ID1812bcc8a02c3b82be8d78402b020b0c]]
= system controller replace show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status of controller replacement

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller replace show` command displays overall information
about the currently running, or previously run controller replacement
operation. The command displays the following information:

* Operation Status
* Error message
* Recommended action



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

== Examples

The following example displays information about automated nondisruptive
operation:
```

cluster1::*> system controller replace show
Node Status Error-Action
--- --- ---
node1 Paused-for-intervention Follow the instructions given in
Step Details

node2 None
Step Details:

Controller replacement operation has been paused for user intervention.
Collect the following info from the current node:
1. vserver services name-service dns show
2. service-processor show -node * -instance
3. network port ifgrp show
4. network port vlan show
5. network interface failover-groups show
6. storage array config show -switch switchname
7. storage encryption disk show
2 entries were displayed.

```
[[IDa1be1a3c5ec3bc0292e11a384e311cca]]
= system controller replace start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start controller replacement

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system controller replace start` command is used to initiate a
controller-replacement operation. The update is preceded by a validation
of the HA pair to ensure that any issues that might affect the update are
identified.

There are predefined points in the update when the update can be paused
(either requested by the user or by the operation in case of an error or
for manual intervention).

== Parameters

``[-nodes <nodename>,...]`` - Nodes for Controller Replacement [.advaced-
priv]#(privilege: advanced)#::
Specifies the nodes that are to be replaced.

``[-simulate <true>]`` - Simulate Controller Replacement [.advaced-
priv]#(privilege: advanced)#::
Dry run the operation. Checks for all validations.
```

```
``[-skip-metrocluster-check {true|false}]`` - Skip Metrocluster Check
before Replacement [.advaced-priv]#(privilege: advanced)#::
Skips the DR mirror offline check when performing MetroCluster validation.
In 4-node MetroCluster configuration, if controller replacement is already
complete one one site, then the partner site should replace its
controllers by setting this parameter to `true` . The default value is
false.

``[-nso {true|false}]`` - Select NSO Procedure [.advaced-priv]#(privilege:
advanced)#::
Follow the negotiated switchover switchback based controller replacement
procedure for 4 node MCC FC system. To select ARL based procedure for 4
node MCC FC , this parameter needs to be set to `false` . This parameter
is only honored in MCC configuration. Therefore, setting or unsetting of
this parameter has no effect in HA configuration. The default value is
true for MCC FC.

== Examples

The following example shows the replacement operation:
```

cluster1::> system controller replace start -nodes node1,node2 -simulate true

Warning: 1. Current version of node is 9.4.0
Before starting controller replacement, please ensure that the new controllers are in the version 9.4.0
2. Verify that NVMEM or NVRAM batteries of the new nodes are charged, and charge them if they are not. You need to physically check the new nodes to see if
the NVMEM or NVRAM batteries are charged. You can check the battery status either by connecting to a console or using SSH, logging into the Service Processor
(SP) for your system, and use the system sensors to see if the battery has a sufficient charge.
Attention: Do not try to clear the NVRAM contents. If there is a need to clear the contents of NVRAM, contact NetApp technical support.
3. If you are replacing the controllers with an used one, please ensure to run wipeconfig before controller replacement
4. Current model name is FAS8040
Before starting the operation, please ensure that the new controller model is supported for controller replacement.

Do you want to continue? {y|n}: y
Controller replacement: Prechecks in progress……
Controller replacement has been paused for user intervention.
Please collect the following info from the current node:
vserver services name-service dns show
network interface show -curr-node node -role cluster,intercluster,node-mgmt,cluster-mgmt
network port show -node node -type physical
service-processor show -node * -instance
network fcp adapter show -node node
network port ifgrp show
network port vlan show
system node show -instance -node node

run -node node sysconfig
storage aggregate show -node node
volume show -node node
network interface failover-groups show
storage array config show -switch switchname
system license show -owner node
storage encryption disk show

```
[[ID227d7672d1ec00b09eb3117b325288c4]]
= system controller service-event delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Manually clear a selected service event

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller service-event delete` command removes the service
event from the list and extinguishes all related FRU attention LEDs.



In some cases, where the underlying fault condition remains, the service
event might be reported again, causing it to reappear in the list. In such
cases, it is necessary to remedy the underlying fault condition in order
to clear the service event.

== Parameters

``-node {<nodename>|local}`` - Node::
Selects service events on the specified nodes.

``-event-id <integer>`` - Service Event ID::
Selects the service events that match the specified event identifier.
Together with the node, this field uniquely identifies the row to delete.
Use the xref:{relative_path}system-controller-service-event-
show.html[system controller service-event show] command to find the event
identifier for the service event to delete.

== Examples

The following example lists the currently active service events. Then,
using the listed Service Event ID, the service event is deleted:
```

cluster1::> system controller service-event show

Node ID Event Location Event Description
--- --- --- ---

plata4-1a 1 DIMM in slot 1 on Controller A Uncorrectable ECC

cluster1::> system controller service-event delete -event-id 1

```
== Related Links

* xref:{relative_path}system-controller-service-event-show.html[system
controller service-event show]




[[IDf62b9c54a3c65c90269d4bb69c0742b1]]
= system controller service-event show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the active service events causing attention LEDs to be lit

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller service-event show` command displays one or more
events that have been detected by the system for which a physical service
action might be required. Physical service actions sometimes involve
replacing or re-seating misbehaving FRUs. In such cases FRU attention LEDs
will be illuminated to assist in physically locating the FRU in need of
attention. When the FRU in question is contained within another FRU, both
the inner and outer FRU attention LEDs will be lit. It creates a path of
LEDs that starts at the chassis level and leads to the FRU in question.
For example, if a DIMM is missing from the controller motherboard, the
storage OS will detect this and log a service event whose location is the
DIMM slot on the controller. The DIMM slot LED, controller LED and chassis
LED will all be lit to create a path of LEDs to follow.



FRU Attention LEDs that are not visible from outside of the system (e.g.
those on the controller motherboard such as DIMMs, boot device etc.) will
remain on for a few minutes, even after power is removed from the
containing FRU. As such, when the controller is removed from the chassis,
a DIMM slot FRU attention LED will remain on, helping to locate the FRU in
```

```
need of attention.



Generally, service events are cleared automatically when the issue is
resolved. The corresponding FRU attention LEDs are extinguished
accordingly. In cases where the service event request is caused by an
environmental issue, it might be necessary to manually remove the service
event from the list. This can be done using the
xref:{relative_path}system-controller-service-event-delete.html[system
controller service-event delete] command.

== 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>|local}]`` - Node::
Selects service events on the specified nodes.

``[-event-id <integer>]`` - Service Event ID::
Selects the service events that match the specified event identifier.
Together with the node, this field uniquely identifies the row for use
with the xref:{relative_path}system-controller-service-event-
delete.html[system controller service-event delete] command

``[-event-loc <text>]`` - Location::
Selects the service events that match the specified event location.

``[-event-desc <text>]`` - Description::
Selects the service events that match the specified event description.

``[-event-timestamp <text>]`` - Timestamp::
The time that the event occurred, recorded by the Service Processor

== Examples

The following example lists the currently active service events.
```

cluster1::> system controller service-event show

Node ID Event Location Event Description
--- --- --- ---
plata4-1a 1 DIMM in slot 1 on Controller A Uncorrectable ECC

```
== Related Links

* xref:{relative_path}system-controller-service-event-delete.html[system
controller service-event delete]




[[ID39ddecf30b7ba165f4d0b26469a980f1]]
= system controller slot module insert
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a module on the controller

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller slot module insert` command adds a module on the
controller.

== Parameters

``-node {<nodename>|local}`` - Node::
Selects the PCIe modules that are present in the specified node.

``-slot <text>`` - Slot Number::
Selects the PCIe modules present in the specified slot or slot-subslot
combination.

== Examples

The following example adds a module in the local node:
```

p2i030::> system controller slot module insert -node local -slot 1

Warning: IO_CARRIER_NIANTIC_NIC module in slot 1 of node p2i030 will be powered

on and initialized.
Do you want to continue? {y|n}:y
The module has been successfully powered on, initialized and placed into service.

p2i030::>

```
[[IDb8d80d4d888e2826eddf16346a27d5b6]]
= system controller slot module remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a module on the controller

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller slot module remove` command removes a module on the
controller.

== Parameters

``-node {<nodename>|local}`` - Node::
Selects the PCIe modules that are present in the specified node.

``-slot <text>`` - Slot Number::
Selects the PCIe modules present in the specified slot or slot-subslot
combination.

== Examples

The following example removes a module in the local node:
```

p2i030::> system controller slot module remove -node local -slot 1

Warning: IO_CARRIER_NIANTIC_NIC module in slot 1 of node p2i030 will be
powered off for removal.
Do you want to continue? {y|n}: y
The module has been successfully removed from service and powered off. It can now be safely removed.

p2i030::>

```
[[ID71bda756d21aa72ba84d5fb190195e1d]]
= system controller slot module replace
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Power off a module on the controller for replacement

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller slot module replace` command powers off a module on
the controller for replacement.

== Parameters

``-node {<nodename>|local}`` - Node::
Selects the PCIe modules that are present in the specified node.

``-slot <text>`` - Slot Number::
Selects the PCIe modules present in the specified slot or slot-subslot
combination.

== Examples

The following example powers off a module in the local node:
```

p2i030::> system controller slot module replace -node local -slot 1

Warning: IO_CARRIER_NIANTIC_NIC module in slot 1 of node p2i030 will be powered
off for replacement.
Do you want to continue? {y|n}: y
The module has been successfully powered off. It can now be safely replaced. After the replacement module is inserted, use the "system controller slot module insert" command to place the module into service.

p2i030::>

```
[[ID0f57272fbd3d03e580dcce16f1ea2c70]]
= system controller slot module show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display hotplug status of a module on the controller

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller slot module show` command displays hotplug status
of a module on the controller. The command displays the following
information about the PCIe modules:

* Node
* Slot
* Module
* Status

To display more details, use the -instance parameter.

== 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>|local}]`` - Node::
Selects the PCIe modules that are present in the specified node.

``[-slot <text>]`` - Slot Number::
Selects the PCIe modules present in the specified slot or slot-subslot
combination.

``[-status <text>]`` - Module Status::
Selects hotplug status for PCIe modules.

``[-card <text>]`` - Module Name::
Selects module name for PCIe modules.

== Examples
```

```
The following example displays hotplug status of PCI modules found in the
local node:
```

::> system controller slot module show -node local
Node Slot Module Status
--- --- --- ---
localhost 1 IO_CARRIER_NIANTIC_NIC powered-on
localhost 2 IO_4X_10GBT_INTL_NIC powered-on
localhost 3 IO_4X_12Gb_PMC_SAS powered-on
localhost 4 IO_4X_10GBE_16GFC_QLGC_CNA powered-on
localhost 5 IO_4X_12Gb_PMC_SAS powered-on
localhost 6 NVRAM10 hotplug-not-supported
localhost 6-1 empty
localhost 6-2 empty
localhost 7 IO_4X_12Gb_PMC_SAS powered-on
localhost 8 IO_4X_10GBT_INTL_NIC powered-on
localhost 9 IO_4X_12Gb_PMC_SAS powered-on
localhost 10 IO_4X_12Gb_PMC_SAS powered-on
localhost 11 IO_4X_12Gb_PMC_SAS powered-on
13 entries were displayed.

::>

```
[[IDee6bfb0b44a44291b9196b73f329a0b7]]
= system controller sp config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Service Processor Config Table

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system controller sp config show` command displays the following
configuration information of the service processor for all nodes in the
cluster:

*  Node name
*  Service processor status
*  Service processor firmware version
*  Booted firmware version
*  Service processor configuration status
*  Physical Ethernet link status of service processor
```

```
*  Health monitor status

To display more details, use the -instance parameter. These commands are
available for 80xx, 25xx and later systems. Earlier models are not
supported.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the field that you specify.

| ``[-instance  ]`` }::
Displays detailed configuration information of the service processor.

``[-node {<nodename>|local}]`` - Node::
Use this parameter to list the service processor configuration of the
specific node.

``[-version <text>]`` - Firmware Version::
Selects the service processor configuration with the specified firmware
version.

``[-boot-version {primary|backup}]`` - Booted Version::
Selects the service processor configuration with the specified version of
the currently booted partition.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the service processor configuration with the specified monitor
name.

``[-sp-status {online|offline|sp-daemon-offline|node-
offline|degraded|rebooting|unknown|updating}]`` - SP Status::
Selects the service processor configuration with the specified status of
service processor.

``[-sp-config {true|false}]`` - Auto Update Configured::
Selects information about the service processor with the specified
configuration status of the service processor.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the service processor configuration information with the specified
service processor status.

``[-link-status {up|down|disabled|unknown}]`` - Public Link Status::
```

```
Selects the service processor configuration with the specified physical
Ethernet link status.

``[-name <text>]`` - Display Name::
Selects the service processor configuration with the specified unique
name.

== Examples

The example below displays configuration of the service processor in all
the nodes in the cluster:
```

cluster1::> system controller sp config show
Firmware Booted Auto Update SP Link
Node Version Version Configured Status Status Status
---- -------- --------- ------------ -------- --------- ------
node1 2.2.2 primary true online up ok
node2 2.2.2 primary true online up ok

```
The example below displays configuration of the service processor of a
particular node in detail:
```

cluster1::> system controller sp config show -instance -node node1
Node: node1
Firmware Version: 2.2.2
Booted Version: primary
Health Monitor Name: controller
SP Status: online
Auto Update Configured: true
Status: ok
Public Link Status: up
Display Name: SP Config

```
[[ID251810ae09d06302698dfc03dc1ad1f2]]
= system controller sp upgrade show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Service Processor Upgrade Table

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `system controller sp upgrade show` command displays the following
information about the service processor firmware of all the nodes in the
cluster:

*  Node name
*  Is new firmware available?
*  Is autoupdate enabled?
*  Status of autoupdate
*  Health monitor status

To display more details, use the -instance parameter. These commands are
available for 80xx, 25xx and later systems. Earlier models are not
supported.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed upgrade information of the service processor.

``[-node {<nodename>|local}]`` - Node::
Use this parameter to list the upgrade information of the service
processor on the specified node.

``[-new-fw-avail {true|false}]`` - New Firmware Available::
Selects the information of the service processors which have new firmware
available.

``[-new-fw-version <text>]`` - New Firmware Version::
Selects the information about service processors with the specified
firmware version.

``[-auto-update {true|false}]`` - Auto Update::
Selects the information about service processors with the specified state.

``[-auto-update-stat {installed|corrupt|updating|auto-updating|none}]`` -
Auto Update Status::
Selects the information about service processors with the specified auto
update status.

``[-auto-update-sttime <MM/DD/YYYY HH:MM:SS>]`` - Auto Update Start Time::
Selects the information about service processors with the specified start
time.
```

```
``[-auto-update-entime <MM/DD/YYYY HH:MM:SS>]`` - Auto Update End Time::
Selects the information about service processors with the specified end
time.

``[-auto-update-per <integer>]`` - Auto Update Percent Done::
Selects the information about service processors with the specified auto
update percentage completed.

``[-auto-update-maxret <integer>]`` - Auto Update Maximum Retries::
Selects the information about service processors with the specified
maximum number of retries.

``[-auto-update-curret <integer>]`` - Auto Update Current Retries::
Selects the information about service processors with the specified number
of current retries.

``[-auto-update-prevstat {failed|passed}]`` - Previous AutoUpdate Status::
Selects the information about service processors with the specified
automatic update status.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor Name::
Selects the information about service processors with the specified
monitor name.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Status::
Selects the information about service processors with the specified health
monitor status.

``[-name <text>]`` - Display Name::
Selects the information about service processors with the specified
display name.

== Examples

The example below displays service processor upgrade information for all
nodes in the cluster:
```

```
cluster1::> system controller sp upgrade show
New Firmware Auto Update Auto Update
Node Available Feature Status Status
---- --------------- ------------ ------------- ---------
node1 false true installed ok
node2 false true installed ok
2 entries were displayed.
```

```
The example below displays the detailed service processor upgrade
information for a specific node:
```

cluster1::> system controller sp upgrade show -instance -node node1
Node: node1
New Firmware Available: false
New Firmware Version: Not Applicable
Auto Update: true
Auto Update Status: installed
Auto Update Start Time: Thu Oct 20 20:06:03 2012 Etc/UTC
Auto Update End Time: Thu Oct 20 20:09:19 2012 Etc/UTC
Auto Update Percent Done: 0
Auto Update Maximum Retries: 5
Auto Update Current Retries: 0
Previous AutoUpdate Status: passed
Health Monitor Name: controller
Status: ok
Display Name: SP Upgrade

```
:leveloffset: -1


= system feature-usage commands

:leveloffset: +1


[[IDb251a39170faed8bef396170687a858d]]
= system feature-usage show-history
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Feature Usage History

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Display feature usage information in the cluster on a per-node and per-
week basis.

== 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>|local}]`` - Node::
Displays feature usage information for the specified node name.

``[-serial-number <Node Serial Number>]`` - Node Serial Number::
Displays feature usage information for the specified serial number.

``[-feature-name <Managed Feature>]`` - Feature Name::
Displays feature usage information for the specified feature name.

``[-week-number <Sequence Number>]`` - Week Number::
Displays feature usage information for the specified week number.

``[-usage-status {not-used|configured|in-use|not-available}]`` - Usage
Status::
Displays feature usage information that matches the specified usage
status.

``[-date-collected <MM/DD/YYYY HH:MM:SS>]`` - Collection Date::
Displays feature usage information that is collected on the day matching
the specified date.

``[-owner <text>]`` - Owner::
Displays feature usage information for the specified owner name.

``[-feature-message <text>]`` - Feature Message::
Displays feature usage information that contains the specified feature
message.

== Examples

The following example displays a usage output filtered by the serial
number and feature name:
```

cluster1::> system feature-usage show-history -serial-number 1-81-0000000000000001122334455 -feature -name NFS
Node Serial Number: 1-81-0000000000000001122334455
Feature Name: NFS
Owner: node1

Week # Usage Status Date Collected Feature Message
------ ---------- ---------- ----------
4 in-use 01/22/13 10:00:00
3 in-use 01/15/13 10:00:00
2 not-used 01/08/13 10:00:00
1 configured 01/01/13 10:00:00

4 entries were displayed.

```
[[IDb5a4158cee436771eff77e3671f5f7d9]]
= system feature-usage show-summary
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Feature Usage Summary

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Display usage summary information about features in the cluster on a per-
node basis. The summary information includes counter information such as
the number of weeks the feature was in use and the last date and time the
feature was used. Additional information can also be displayed by using
the -instance parameter.

== 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.

``[-serial-number <Node Serial Number>]`` - Node Serial Number::
Displays usage summary information for the specified serial number.

``[-feature-name <Managed Feature>]`` - Feature Name::
Displays usage summary information for the specified feature name.
```

```
``[-weeks-in-use <integer>]`` - Weeks In-Use::
Displays usage summary information for features matching the number of
weeks in use.

``[-last-used <MM/DD/YYYY HH:MM:SS>]`` - Date last used::
Displays usage summary information for features last used on the specified
date.

``[-owner <text>]`` - Owner::
Displays usage summary information for the specified owner name.

``[-weeks-not-used <integer>]`` - Weeks Not Used::
Displays usage summary information for features matching the number of
weeks not in use.

``[-weeks-configured <integer>]`` - Weeks Configured::
Displays usage summary information for features matching the number of
weeks that the feature was in configuration.

``[-weeks-not-available <integer>]`` - Weeks Data Not Available::
Displays usage summary information for features matching the number of
weeks when usage data was not available.

== Examples

The following example displays a usage summary output for a cluster of two
nodes:
```

cluster1::> system feature-usage show-summary
Node Serial Number: 1-81-0000000000000000001122334455
Owner: node1
Feature Name Weeks In Use Date Last Used
---------- ---------- ----------
CIFS 10 1/1/2013 23:27:49
NFS 15 1/8/2013 23:48:03

Node Serial Number: 1-81-0000000000000000001122334466
Owner: node2
Feature Name Weeks In Use Date Last Used
---------- ---------- ----------
CIFS 10 1/1/2013 23:26:38
NFS 20 1/8/2013 23:46:48
4 entries were displayed.

```
:leveloffset: -1



= system fru-check commands
```

```
:leveloffset: +1


[[ID1d4a8bc8247881b776ecd9ba651e401b]]
= system fru-check show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Information About the FRUs in the Controller

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system fru-check show` command checks and displays the results of
quick diagnostic tests done for certain FRUs of each controller in the
cluster. The tests are not intended to be exhaustive, but simply to do a
quick check of certain FRUs especially after replacement.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that have the specified name.

| ``[-instance  ]`` }::
Selects detailed information (if available) for all the FRUs.

``[-node {<nodename>|local}]`` - Node::
Selects the FRUs that belong to the node that has the specified name.

``[-serial-number <text>]`` - FRU Serial Number::
Selects the FRU matching the specified serial number.

``[-fru-name <text>]`` - FRU Name::
Selects the FRU matching the specified fru-name.

``[-fru-type {controller|dimm|bootmedia|nvram|nvdimm}]`` - FRU Type::
Selects the FRUs of the specified type.

``[-fru-status {pass|fail|unknown}]`` - Status::
Selects the FRUs whose FRU check status matches that specified. "pass"
indicates the FRU is operational. "fail" indicates the FRU is not
```

```
operating correctly. "unknown" indicates a failure to obtain FRU
information during the check.

``[-display-name <text>]`` - Display Name::
Selects the FRU matching the specified display name.

``[-location <text>]`` - Location::
Selects the FRUs whose location matches that specified. Example: Memory
Slot: 1

``[-additional-info <text>]`` - Additional Info::
Selects the FRUs whose additional information matches that specified.
Example: Part No: 69003140-I00-NTA-T

``[-reason <text>]`` - Details::
Selects the FRUs whose failure reason matches that specified.




:leveloffset: -1


= system ha commands

:leveloffset: +1


[[ID314c375a8d9621f0deb8eaa6d37b3f7a]]
= system ha interconnect config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the high-availability interconnect configuration information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect config show` command displays the high-
availability interconnect device basic configuration information.

== Parameters
```

```
``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display all the fields from all nodes in cluster.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display all the fields from the specified node in
the cluster.

``[-transport <text>]`` - Interconnect Type [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this HA interconnect transport type.

``[-local-sysid <integer>]`` - Local System ID [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this local system unique identifier.

``[-partner-sysid <integer>]`` - Partner System ID [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this partner system unique identifier.

``[-initiator {local|partner}]`` - Connection Initiator [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
initiator of the connection request.

``[-port-name <text>,...]`` - Port [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this port name.

``[-ipaddress <text>,...]`` - IP Address [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this IP address.

``[-interface {backplane|external}]`` - Interface [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. ``_external_`` means
the HA interconnect links between partner nodes are connected externally.
``_backplane_`` means the HA interconnect links between partner nodes are
connected over the backplane.

== Examples
```

```
The following example displays the HA interconnect configuration
information on FAS8000 series nodes in the cluster:
```

cluster1::*> system ha interconnect config show
Node: ic-f8040-01
Interconnect Type: Infiniband (Mellanox ConnectX)
Local System ID: 536875713
Partner System ID: 536875678
Connection Initiator: local
Interface: backplane

Port IP Address Flags
---- ---------- ----------
ib0a 192.0.3.236 0x0
ib0b 192.0.3.237 0x0
Node: ic-f8040-02
Interconnect Type: Infiniband (Mellanox ConnectX)
Local System ID: 536875678
Partner System ID: 536875713
Connection Initiator: partner
Interface: backplane

Port IP Address Flags
---- ---------- ----------
ib0a 192.0.3.96 0x0
ib0b 192.0.3.97 0x0

2 entries were displayed.

```
The following example displays the HA interconnect configuration
information on FAS2500 series nodes in the cluster:
```

cluster1::*> system ha interconnect config show
Node: ic-f2554-03
Interconnect Type: Infiniband (Mellanox Sinai)
Local System ID: 1781036608
Partner System ID: 1780360209
Connection Initiator: local
Interface: backplane

Port IP Address Flags
---- ---------- ----------
ib0a ib0a -
Node: ic-f2554-04
Interconnect Type: Infiniband (Mellanox Sinai)
Local System ID: 1780360209
Partner System ID: 1781036608
Connection Initiator: partner
Interface: backplane

Port IP Address Flags

---- ---------- ----------
ib0a ib0a -

2 entries were displayed.

```
[[ID52d9595beea7f52f037cb528d4fda6d7]]
= system ha interconnect link off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Turn off the interconnect link

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect link off` command turns off the specified link
on the high-availability interconnect device. For the nodes in the cluster
with two external high-availability interconnect links, you must specify
the link number (0-based) to turn off the specified link. For the nodes in
the cluster with interconnect links over the backplane, you must specify
the link number 1 to turn off the link.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies the node on which the interconnect link
is to be turned off. The value "local" specifies the current node.

``-link {0|1}`` - Link [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies the interconnect link number (0-based)
to turn off.

== Examples

The following example displays output of the command on the nodes with a
single interconnect link or nodes with interconnect links over the
backplane:
```

cluster1::*> system ha interconnect link off -node ic-f3250-02 -link 0

Error: command failed: Invalid link value 0. Specify 1.

cluster1::*> system ha interconnect link off -node ic-f3250-02 -link 1

```
The following example displays output of the command on the nodes with two
interconnect links connected externally:
```

cluster1::*> system ha interconnect link off -node ic-f3250-02 -link 0

cluster1::*> system ha interconnect link off -node ic-f3250-02 -link 1

```
[[ID56b1dbc6dd86dad50e82580581ea6869]]
= system ha interconnect link on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Turn on the interconnect link

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect link on` command turns on the specified link
on the high-availability interconnect device. For the nodes in the cluster
with two external high-availability interconnect links, you must specify
the link number (0-based) to turn on the specified link. For the nodes in
the cluster with interconnect links over the backplane, you must specify
the link number 1 to turn on the link.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies the node on which the interconnect link
is to be turned on. The value "local" specifies the current node.

``-link {0|1}`` - Link [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies the interconnect link number (0-based)
to turn on.

== Examples

The following example displays output of the command on the nodes with a
single interconnect link or nodes with interconnect links over the
backplane:
```

cluster1::*> system ha interconnect link on -node ic-f3250-02 -link 0

Error: command failed: Invalid link value 0. Specify 1.

cluster1::*> system ha interconnect link on -node ic-f3250-02 -link 1

```
The following example displays output of the command on the nodes with two
interconnect links connected externally:
```

cluster1::*> system ha interconnect link on -node ic-f3250-02 -link 0

cluster1::*> system ha interconnect link on -node ic-f3250-02 -link 1

```
[[ID259178e8a1f9048adffadbd6a6ddd146]]
= system ha interconnect ood clear-error-statistics
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear error statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood clear-error-statistics` command enables
you to clear all the error statistics collected for the out-of-order
delivery-capable high-availability interconnect device. This command is
only supported on FAS2500 series nodes in the cluster.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies which node will have the error
statistics cleared. The value "local" specifies the current node.

== Examples
```

cluster1::*> system ha interconnect ood clear-error-statistics -node ic-f2554-03

```
[[ID01a61d2cb388d0b1649bf1105fb9acf5]]
= system ha interconnect ood clear-performance-statistics
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear performance statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood clear-performance-statistics` command
enables you to clear all the performance statistics collected for the out-
of-order delivery-capable high-availability interconnect device. This
command is only supported on FAS2500 series nodes in the cluster.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies which node will have the performance
statistics cleared. The value "local" specifies the current node.

== Examples
```

cluster1::*> system ha interconnect ood clear-performance-statistics -node ic-f2554-03

```
[[IDa59998433eb634ef46f8b740a6672555]]
= system ha interconnect ood disable-optimization
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable coalescing work requests

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood disable-optimization` command disables the
optimization capability on the high-availability interconnect device. The
command is only supported on FAS2500 series nodes in the cluster.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies which node will have the optimization
disabled. The value "local" specifies the current node.

== Examples
```

cluster1::*> system ha interconnect ood disable-optimization -node ic-f2554-03

```
[[IDa309a84a328c7eb524c2dcd8fbe08ff8]]
= system ha interconnect ood disable-statistics
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable detailed statistics collection

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood disable-statistics` command disables
collection of the statistics on the out-of-order delivery-capable high-
availability interconnect device. This command is only supported on
FAS2500 series nodes in the cluster.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies which node will have the statistics
collection disabled. The value "local" specifies the current node.

== Examples
```

cluster1::*> system ha interconnect ood disable-statistics -node ic-f2554-03

```
[[ID752a2f991c20a5b7295194262cf64033]]
= system ha interconnect ood enable-optimization
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable coalescing work requests

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood enable-optimization` command enables you
to turn on optimization (coalescing out-of-order delivery requests) on the
high-availability interconnect device. This command is only supported on
FAS2500 series nodes in the cluster.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies which node will have the optimization
enabled. The value "local" specifies the current node.

== Examples
```

cluster1::*> system ha interconnect ood enable-optimization -node ic-f2554-03

```
[[IDc487624c65f60a6b574deffaae20810f]]
= system ha interconnect ood enable-statistics
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable detailed statistics collection

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood enable-statistics` command enables
collection of the statistics on the out-of-order delivery-capable high-
availability interconnect device. This command is only supported on
FAS2500 series nodes in the cluster.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies which node will have the statistics
collection enabled. The value "local" specifies the current node.

== Examples
```

cluster1::*> system ha interconnect ood enable-statistics -node ic-f2554-03

```
[[ID3fb0a1c63f89348b5fa104da94d60b98]]
= system ha interconnect ood send-diagnostic-buffer
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Send diagnostic buffer to partner

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood send-diagnostic-buffer` command enables
you to run a short out-of-order delivery diagnostic test. The command
sends a buffer to the partner controller over the high-availability
interconnect. This command is only supported on FAS2500 series nodes in
the cluster.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This mandatory parameter specifies which node will send the diagnostic
buffer to its partner. The value "local" specifies the current node.

== Examples

The following example demonstrates how to use this command to send a
diagnostic buffer to the partner:
```

cluster1::*> system ha interconnect ood send-diagnostic-buffer -node ic-f2554-03

```
[[ID5029a9656a2db688da3cb004a80584b7]]
= system ha interconnect ood status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the high-availability interconnect device out-of-order delivery
(OOD) information
```

```
*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect ood status show` command displays
configuration information of the out-of-order delivery-capable high-
availability interconnect devices. This command is supported only on
FAS2500 series nodes in the cluster.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display all the fields from all nodes in cluster.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display all the fields from the specified node in
the cluster.

``[-is-ood-enabled {true|false}]`` - Is OOD Enabled [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-is-coalescing-enabled {true|false}]`` - Is Coalescing Enabled
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

== Examples

The following example displays the HA interconnect device out-of-order
delivery configuration information on FAS2500 series nodes in the cluster.
```

cluster1::*> system ha interconnect ood status show
Node: ic-f2554-03
NIC Used: 0
Is OOD Enabled: true
Is Coalescing Enabled: true
Node: ic-f2554-04
NIC Used: 0
Is OOD Enabled: true
Is Coalescing Enabled: true
2 entries were displayed.

```
[[ID5d06670368acd821a1d2804afdabfdf6]]
= system ha interconnect port show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the high-availability interconnect device port information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect port show` command displays the high-
availability interconnect device port physical layer and link layer status
information.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display all the fields from all nodes in the
cluster.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display all the fields from the specified node in
the cluster.

``[-link-monitor {on|off}]`` - Link Monitor Detection [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-port <integer>,...]`` - Port Number [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value.

``[-phy-layer-state {invalid|sleep|polling|disabled|port-configuration-
testing|linkup|link-error-recovery|phytest|reserved}]`` - Physical Layer
State [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.
```

```
``[-link-layer-state {invalid|down|initialize|armed|active|reserved}]`` -
Link Layer State [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-phy-link-up-count <integer>,...]`` - Physical Link Up Count [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is total
number of times the link on a given port is transitioned up.

``[-phy-link-down-count <integer>,...]`` - Physical Link Down Count
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is total
number of times the link on a given port is transitioned down.

``[-is-active-link {true|false}]`` - Is the Link Active [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value ``_true_``
means the interconnect data channels are established on this link.

== Examples

The following example displays the HA interconnect device port information
on FAS8000 series nodes in the cluster:
```

cluster1::*> system ha interconnect port show
Physical Link
Link Layer Layer Physical Physical Active
Node Monitor Port State State Link Up Link Down Link
--- --- --- --- --- --- --- ---
ic-f8040-01 on
0 linkup active 1 0 true
1 linkup active 1 0 false
ic-f8040-02 on
0 linkup active 1 0 true
1 linkup active 1 0 false
2 entries were displayed.

```
[[ID30cd968a4d3d3f36549f54f5bd9154f3]]
= system ha interconnect port sharing modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]


*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect port sharing modify` command sets the state of
port sharing for the specified node.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the node on which to modify port sharing.

``[-enabled {true|false}]`` - Enabled [.advaced-priv]#(privilege:
advanced)#::
Set this parameter to enable/disable port sharing.

== Examples

The following example modifies the HA interconnect device port sharing to
enabled:
```

cluster1::*> system ha interconnect port sharing modify -node node1 -enabled true

cluster1::*> system ha interconnect port sharing show
Node Enabled
--- ---
node1 true
node2 false

```
[[ID07af926a0d502261a3721874860bb326]]
= system ha interconnect port sharing show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the high-availability interconnect device port sharing information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect port sharing show` command displays the high-
availability interconnect device port sharing status information.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display all the fields for each node.

``[-node <nodename>]`` - Node [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information from the specified node in the
cluster.

``[-enabled {true|false}]`` - Enabled [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value.

== Examples

The following example displays the HA interconnect device port sharing
information:
```

cluster1::*> system ha interconnect port sharing show
Node Enabled
---------- -------
node1 true
node2 false

```
[[ID10380322cc36ac96a173c4c99d204222]]
= system ha interconnect statistics clear-port-symbol-error
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear the high-availability interconnect port symbol errors

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect statistics clear-port-symbol-error` command
clears the high-availability interconnect device port symbol errors. This
command is supported only on FAS2500 series nodes in the cluster.

NOTE: To display the high-availability interconnect device port
statistics, use the xref:{relative_path}statistics-show.html[statistics
show -object ic_hw_port_stats] command.



== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

== Examples
```

cluster1::*> system ha interconnect statistics clear-port-symbol-error -node ic-f2554-03

```
== Related Links

* xref:{relative_path}statistics-show.html[statistics show]




[[IDed7d547817042f15d87debed258b40ce]]
= system ha interconnect statistics clear-port
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear the high-availability interconnect port counters

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect statistics clear-port` command clears the
high-availability interconnect device port statistics. This command is
supported only on FAS2500 series and FAS8000 series nodes in the cluster.

NOTE: To display the high-availability interconnect device port
statistics, use the xref:{relative_path}statistics-show.html[statistics
show -object ic_hw_port_stats] command.



== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

== Examples
```

cluster1::*> system ha interconnect statistics clear-port -node ic-f8040-01

```
== Related Links

* xref:{relative_path}statistics-show.html[statistics show]
```

```
[[IDc1c33569f15537a1532672893b6f8aff]]
= system ha interconnect statistics show-scatter-gather-list
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the high-availability interconnect scatter-gather list entry
statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect statistics show-scatter-gather-list` command
displays the high-availability interconnect device scatter-gather list
entry statistics. Out of all possible 32 entries in a scatter-gather list,
the command displays only the entries that have valid data.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display all the fields from all nodes in cluster.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display all the fields from the specified node in
the cluster.

``[-sge <integer>,...]`` - Scatter-Gather Entry [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this scatter-gather element index value.

``[-total-count <integer>,...]`` - Total Count [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of times a particular scatter-gather list element is used.
```

```
``[-total-size <integer>,...]`` - Total Size [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of bytes written by the high-availability interconnect device using
a particular scatter-gather list element.

== Examples
```

cluster1::*> system ha interconnect statistics show-scatter-gather-list
Node: ic-f8040-01
Entry Count Size
----- ---------- ----------
1 410925 77344493
2 988 1246987
3 72 747325
4 93264 1527155579
8 9 294912
9 9 294912

Node: ic-f8040-02
Entry Count Size
----- ---------- ----------
1 1544405 310004390
2 6217 16779908
3 1222 12003411
4 338606 5543436659
6 2 41980
7 2 46136
8 18 589824
9 18 589824

2 entries were displayed.

```
[[ID0c1cd6996729310fa5de116af86ba68a]]
= system ha interconnect statistics performance show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the high-availability interconnect device performance statistics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system ha interconnect statistics performance show` command displays
```

```
the high-availability interconnect device performance statistics.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display all the fields from all nodes in cluster.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display all the fields from the specified node in
the cluster.

``[-elapsed <integer>]`` - Elapsed Time (secs) [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. Displays the total
elapsed time between statistics collection start time to end time. During
the initialization stage, statistics collection starts when the partner
node is up and ready. After the initialization stage, the statistics
collection start time is reset after every execution of this command. This
means that after the initialization stage, elapsed time represents the
time between current command execution and previous command execution.

``[-qmax-wait <integer>]`` - Maximum Queue Wait Count [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this wait value. The queue maximum wait value
is the total number of times the interconnect device waited to post
requests on the send queue.

``[-qmax-wait-time <integer>]`` - Average Queue Wait Time (usecs)
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this average wait time value. The queue
maximum wait time is the average amount of time the interconnect device
waited to post requests on the send queue.

``[-qmax-timeout <integer>]`` - Maximum Queue Timeouts [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The queue maximum
timeout value is the total number of times the interconnect device timed
out waiting to post requests on the send queue.

``[-preempt-timeout <integer>]`` - Preempt Timeouts [.advaced-
priv]#(privilege: advanced)#::
```

```
Selects the nodes that match this parameter value. The timeout value is
the total number of times polling on the given transfer ID is preempted.

``[-nonpreempt-timeout <integer>]`` - Non-Preempt Timeouts [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The timeout value is
the total number of times polling on the given transfer ID stopped due to
interconnect device read/write timeout.

``[-notify-timeout <integer>]`` - Notify Timeouts [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The timeout value is
the total number of times data transfer on the HA interconnect timed out.

``[-avg-rnv-msgs-time <integer>]`` - Remote NV Messages Average Time
(usecs) [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average time between remote NV messages.

``[-rnv-transfers <integer>]`` - Total Remote NV Transfers [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of remote NV transfers attempted.

``[-avg-rnv-transfer-size <integer>]`` - Remote NV Average Transfer Size
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average remote NV message transfer size.

``[-avg-rnv-transfer-time <integer>]`` - Remote NV Transfers Average Time
(usecs) [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average transfer time taken by remote NV messages.

``[-ic-waits <integer>]`` - Total Count of IC waits for Given ID
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of times the interconnect device waits until the transfer of a
given ID is successful.

``[-ic-waitdone-time <integer>]`` - Average IC Waitdone Time (usecs)
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average time the interconnect device spent waiting for the IDs to be
transferred successfully.

``[-ic-isdone <integer>]`` - Total IC isdone Checks [.advaced-
```

```
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of times the interconnect client checked for the completion of a
given transfer ID.

``[-ic-isdone-pass <integer>]`` - Total IC isdone Checks Success
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of times the check for the completion of a given transfer ID is
successful.

``[-ic-isdone-fail <integer>]`` - Total IC isdone Checks Failed [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of times the check for the completion of a given transfer ID is not
successful.

``[-ic-small-writes <integer>]`` - IC Small Writes [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of <4K size writes performed by the interconnect device.

``[-ic-4k-writes <integer>]`` - IC 4K Writes [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of 4K size writes performed by the interconnect device.

``[-ic-8k-writes <integer>]`` - IC 8K Writes [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of 8K size writes performed by the interconnect device.

``[-ic-16k-writes <integer>]`` - IC 16K+ Writes [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of 16K or more size writes performed by the interconnect device.

``[-ic-xorder-writes <integer>]`` - IC XORDER Writes [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of out-of-order writes performed by the interconnect device.

``[-ic-xorder-reads <integer>]`` - IC XORDER Reads [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of out-of-order reads performed by the interconnect device.
```

```
``[-rdma-read <integer>]`` - RDMA Reads Count [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of RDMA reads performed by the interconnect device.

``[-rdma-read-waitdone-time <integer>]`` - Average IC Waitdone RDMA-READ
Time (usecs) [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average time the interconnect device spent polling for transfer IDs on the
RDMA-read channel.

``[-avg-mbytes-second <text>]`` - Average MegaBytes Transferred per second
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average megabytes (MB) transferred per second.

``[-avg-bytes-transfer <integer>]`` - Average Bytes per Transfer
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average amount of bytes sent per transfer.

``[-total-transfers <integer>]`` - Total Transfers [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of transfers made through the interconnect device.

``[-avg-nvlog-sync-time <integer>]`` - Average Time for NVLOG Sync (msecs)
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average time taken to sync NVLOG between HA partner nodes.

``[-max-nvlog-sync-time <integer>]`` - Maximum Time for NVLOG Sync (msecs)
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
maximum time taken to sync NVLOG between HA partner nodes.

``[-max-sgl-length <integer>]`` - Maximum Scatter-Gather Elements in a
List [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
maximum length of the scatter-gather list supported by the interconnect
device.

``[-ic-recq-waits <integer>]`` - Total Receive Queue Waits to Post Buffer
[.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the total
number of times the interconnect device waited to post an empty buffer
```

```
into the receive queue.

``[-avg-recq-wait-time <integer>]`` - Average Time Receive Queue Waited
(usecs) [.advaced-priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value is the
average amount of time the interconnect device waited to post an empty
buffer into the receive queue.

== Examples

The following example displays the HA interconnect device performance
statistics for FAS8000 series nodes in the cluster:
```

cluster1::*> system ha interconnect statistics performance show
Node: ic-f8040-01
Elapsed Time (secs): 6
Maximum Queue Wait Count: 33
Average Queue Wait Time (usecs): 30
Remote NV Messages Average Time (usecs): 1437
Total Remote NV Transfers: 9297
Remote NV Average Transfer Size: 348
Remote NV Transfers Average Time (usecs): 680
Total IC waits for Given ID: 159
Average IC Waitdone Time (usecs): 5
Total IC isdone Checks: 608
Total IC isdone Checks Success: 608
Total IC isdone Checks Failed: 0
IC Small Writes: 10129
IC 4K Writes: 10
IC 8K Writes: 54
IC 16K+ Writes: 92
IC XORDER Writes: 4855
IC XORDER Reads: 0
RDMA Read Count: 172
Average IC Waitdone RDMA-READ Time (usecs): 0
Average MB/s: 0.98114
Average Bytes per Transfer: 180
Total Transfers: 20720
Average Time for NVLOG Sync (msecs): 1409
Maximum Time for NVLOG Sync (msecs): 1409
Maximum Scatter-Gather Elements in a List: 32
Total Receive Queue Waits to Post Buffer: 0
Node: ic-f8040-02
Elapsed Time (secs): 12
Maximum Queue Wait Count: 29
Average Queue Wait Time (usecs): 68
Remote NV Messages Average Time (usecs): 1386
Total Remote NV Transfers: 19190
Remote NV Average Transfer Size: 375
Remote NV Transfers Average Time (usecs): 670
Total IC waits for Given ID: 304

Average IC Waitdone Time (usecs): 5
Total IC isdone Checks: 1409
Total IC isdone Checks Success: 1409
Total IC isdone Checks Failed: 0
IC Small Writes: 20964
IC 4K Writes: 5
IC 8K Writes: 99
IC 16K+ Writes: 229
IC XORDER Writes: 10261
IC XORDER Reads: 0
RDMA Read Count: 337
Average IC Waitdone RDMA-READ Time (usecs): 0
Average MB/s: 0.57080
Average Bytes per Transfer: 187
Total Transfers: 42883
Average Time for NVLOG Sync (msecs): 1009
Maximum Time for NVLOG Sync (msecs): 1009
Maximum Scatter-Gather Elements in a List: 32
Total Receive Queue Waits to Post Buffer: 0

2 entries were displayed

```
The following example displays the HA interconnect device performance
statistics for FAS2500 series nodes in the cluster:
```

cluster1::*> system ha interconnect statistics performance show
Node: ic-f2554-03
Elapsed Time (secs): 253
Maximum Queue Wait Count: 11
Average Queue Wait Time (usecs): 6837
Maximum Queue Timeouts: 0
Preempt Timeouts: 0
Non-Preempt Timeouts: 0
Notify Timeouts: 0
Remote NV Messages Average Time (usecs): 3343
Total Remote NV Transfers: 59643
Remote NV Average Transfer Size: 8715
Remote NV Transfers Average Time (usecs): 4258
Total IC waits for Given ID: 180
Average IC Waitdone Time (usecs): 3187
Total IC isdone Checks: 499981
Total IC isdone Checks Success: 59922
Total IC isdone Checks Failed: 440059
IC Small Writes: 98722
IC 4K Writes: 5747
IC 8K Writes: 7719
IC 16K+ Writes: 25793
IC XORDER Writes: 66735
IC XORDER Reads: 0
RDMA Read Count: 574
Average IC Waitdone RDMA-READ Time (usecs): 229
Average MB/s: 2.1207

Average Bytes per Transfer: 4680
Total Transfers: 138302
Average Time for NVLOG Sync (msecs): 1236
Maximum Time for NVLOG Sync (msecs): 1236
Maximum Scatter-Gather Elements in a List: 27
Node: ic-f2554-04
Elapsed Time (secs): 257
Maximum Queue Wait Count: 7
Average Queue Wait Time (usecs): 10172
Maximum Queue Timeouts: 0
Preempt Timeouts: 0
Non-Preempt Timeouts: 0
Notify Timeouts: 0
Remote NV Messages Average Time (usecs): 4237
Total Remote NV Transfers: 47134
Remote NV Average Transfer Size: 9559
Remote NV Transfers Average Time (usecs): 5463
Total IC waits for Given ID: 178
Average IC Waitdone Time (usecs): 1890
Total IC isdone Checks: 393191
Total IC isdone Checks Success: 47382
Total IC isdone Checks Failed: 345809
IC Small Writes: 78369
IC 4K Writes: 3815
IC 8K Writes: 6005
IC 16K+ Writes: 22993
IC XORDER Writes: 53529
IC XORDER Reads: 0
RDMA Read Count: 524
Average IC Waitdone RDMA-READ Time (usecs): 62
Average MB/s: 2.3682
Average Bytes per Transfer: 5143
Total Transfers: 111501
Average Time for NVLOG Sync (msecs): 822
Maximum Time for NVLOG Sync (msecs): 822
Maximum Scatter-Gather Elements in a List: 27

2 entries were displayed.

```
[[ID76d7716fd0c3c13e1e415005b9c9c1e4]]
= system ha interconnect status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the high-availability interconnect connection status

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.
```

```
== Description

The `system ha interconnect status show` command displays the high-
availability interconnect connection status. Connection status information
displayed by this command varies by controller model. For nodes with two
HA interconnect links over the backplane or connected externally, this
command displays the following information:

* Node
* Link status on the first port
* Link status on the second port
* Is the link on first port active?
* Is the link on second port active?
* Interconnect RDMA status

For nodes with a single HA interconnect link, this command displays
following the information:

* Node
* Link status
* Interconnect RDMA status



Running the command with the `-instance` or `-node` parameter displays
detailed information about the interconnect device and its ports.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display all the fields for the specified node or all
the nodes.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display all the fields for the specified node.

``[-link-status {up|down}]`` - Link Status [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value ``_up_``
means link is online.

``[-link0-status {up|down}]`` - Link 0 Status [.advaced-priv]#(privilege:
```

```
advanced)#::
Selects the nodes that match this parameter value. The value ``_up_``
means link is online.

``[-link1-status {up|down}]`` - Link 1 Status [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value ``_up_``
means link is online.

``[-ic-rdma {up|down}]`` - IC RDMA Connection [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value. The value ``_up_``
means active interconnect connection with its partner.

``[-is-link0-active {true|false}]`` - Is Link 0 Active [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value ``_true_``
means the interconnect data channels are established on this link.

``[-is-link1-active {true|false}]`` - Is Link 1 Active [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. The value ``_true_``
means the interconnect data channels are established on this link.

``[-slot <integer>]`` - Slot Number [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this PCI slot number.

``[-driver-name <text>]`` - Driver Name [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this interconnect device driver name.

``[-firmware <text>]`` - Firmware Revision [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this firmware version.

``[-version <text>]`` - Version Number [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value.

``[-device-type <text>]`` - Device Type [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this interconnect device type.

``[-serial-number <text>]`` - Serial Number [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this interconnect device serial number.
```

```
``[-debug-firmware {yes|no}]`` - Debug Firmware [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-command-revision <integer>]`` - Command Revision [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this interconnect device command revision.

``[-hardware-revision <integer>]`` - Hardware Revision [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this interconnect device hardware revision.

``[-port1 <integer>]`` - Port Number 1 [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value.

``[-port1-port-name <text>]`` - Port Name [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this port name.

``[-port1-gid <text>]`` - Global Identifier [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this global identifier value.

``[-port1-base-lid <text>]`` - Base Local Identifier [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this base local identifier value.

``[-port1-rm-lid <text>]`` - Remote Local Identifier [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this remote local identifier value.

``[-port1-mtu <integer>]`` - Maximum Transmission Unit [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-port1-data-rate <text>]`` - Data Rate [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value.

``[-port1-link-info <text>]`` - Link Information [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-port1-qsfp-vendor <text>]`` - QSFP Vendor [.advaced-priv]#(privilege:
advanced)#::
```

```
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
vendor name.

``[-port1-qsfp-part-number <text>]`` - QSFP Part Number [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
part-number.

``[-port1-qsfp-type <text>]`` - QSFP Type [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
type.

``[-port1-qsfp-serial-number <text>]`` - QSFP Serial Number [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
serial number.

``[-port2 <integer>]`` - Port Number 2 [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value.

``[-port2-port-name <text>]`` - Port Name [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this port name.

``[-port2-gid <text>]`` - Global Identifier [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this global identifier value.

``[-port2-base-lid <text>]`` - Base Local Identifier [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this base local identifier value.

``[-port2-rm-lid <text>]`` - Remote Local Identifier [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this remote local identifier value.

``[-port2-mtu <integer>]`` - Maximum Transmission Unit [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-port2-data-rate <text>]`` - Data Rate [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this parameter value.

``[-port2-link-info <text>]`` - Link Information [.advaced-
```

```
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value.

``[-port2-qsfp-vendor <text>]`` - QSFP Vendor [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
vendor name.

``[-port2-qsfp-part-number <text>]`` - QSFP Part Number [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
part number.

``[-port2-qsfp-type <text>]`` - QSFP Type [.advaced-priv]#(privilege:
advanced)#::
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
type.

``[-port2-qsfp-serial-number <text>]`` - QSFP Serial Number [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable)
serial number.

== Examples

The following example displays status information about the HA
interconnect connection on FAS8000 series nodes with two HA interconnect
links in the cluster:
```

cluster1::*> system ha interconnect status show
Node: ic-f8040-01
Link 0 Status: up
Link 1 Status: up
Is Link 0 Active: true
Is Link 1 Active: false
IC RDMA Connection: up
Node: ic-f8040-02
Link 0 Status: up
Link 1 Status: up
Is Link 0 Active: true
Is Link 1 Active: false
IC RDMA Connection: up
2 entries were displayed.

```
The following example displays status information about the HA
interconnect connection on FAS2500 series nodes with a single HA
interconnect link in the cluster:
```

cluster1::*> system ha interconnect status show
Node: ic-f2554-01
Link Status: up
IC RDMA Connection: up
Node: ic-f2554-02
Link Status: up
IC RDMA Connection: up
2 entries were displayed.

```
The following example displays detailed information about the HA
interconnect link when parameters like `-instance` , `-node` are used with
the `system ha interconnect status show ` command
```

cluster1::*> system ha interconnect status show -instance -node ic-f8040-01
Node: ic-f8040-01
Link 0 Status: up
Link 1 Status: up
Is Link 0 Active: true
Is Link 1 Active: false
IC RDMA Connection: up
Slot: 0
Driver Name: IB Host Adapter i0 (Mellanox ConnectX MT27518 rev. 0)
Firmware: 2.11.534
Debug Firmware: no

Interconnect Port 0 :
Port Name: ib0a
GID: fe80:0000:0000:0000:00a0:9800:0030:33ec
Base LID: 0x3ec
MTU: 4096
Data Rate: 40 Gb/s (4X) QDR
Link Information: ACTIVE

Interconnect Port 1 :
Port Name: ib0b
GID: fe80:0000:0000:0000:00a0:9800:0030:33ed
Base LID: 0x3ed
MTU: 4096
Data Rate: 40 Gb/s (4X) QDR
Link Information: ACTIVE

```
:leveloffset: -1


= system health commands

:leveloffset: +1


[[ID23f1ffe6c30acacece9a3844fb179c90]]
```

```
= system health alert delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete system health alert

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health alert delete` command deletes all the alerts on the
cluster with the specified input parameters.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to delete alerts generated on a cluster only on the
node you specify.

``-monitor <hm_type>`` - Monitor::
Use this parameter to delete alerts generated on a cluster only on the
monitor you specify.

``-alert-id <text>`` - Alert ID::
Use this parameter to delete alerts generated on a cluster only on the
alert ID you specify.

``-alerting-resource <text>`` - Alerting Resource::
Use this parameter to delete alerts generated on a cluster on the alerting
resource you specify.

== Examples

This example shows how to delete an alert with the specified alert-id:
```

cluster1::> system health alert delete -alert-id DualPathToDiskShelf_Alert -alerting-resource *

```
[[ID022bb7a3df802530060c0183ce3df53e]]
= system health alert modify
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify system health alert

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health alert modify` command suppresses alerts generated on
the cluster and sets the acknowledgement state for an alert.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node on which you want to change the
state.

``-monitor <hm_type>`` - Monitor::
Use this parameter to specify the monitor name on which you want to change
the state.

``-alert-id <text>`` - Alert ID::
Use this parameter to specify the alert ID on which you want to change the
state.

``-alerting-resource <text>`` - Alerting Resource::
Use this parameter to specify the alerting resource name on which you want
to change the state.

``[-acknowledge {true|false}]`` - Acknowledge::
Use this parameter to set the acknowledgement state to true or false.

``[-suppress {true|false}]`` - Suppress::
Use this parameter to set the suppress state to true or false.

``[-acknowledger <text>]`` - Acknowledger::
Use this parameter to set the acknowledger as the filter for setting
state.

``[-suppressor <text>]`` - Suppressor::
Use this parameter to set the suppressor as the filter for setting state.

== Examples

This example modifies the alert field states on the cluster:
```

```
cluster1::> system health alert modify -node * -alert-id DualPathToDiskShelf_Alert -suppress true
```

```
[[ID0b49a6859da4e8508df7d9c699338e5b]]
= system health alert show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
View system health alerts

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health alert show` command displays information about all the
alerts generated on the system. Using -instance will add detailed
information.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays the following additional information about each alert:
+
  * Node name
  * Resource name
  * Severity of the alert
  * Time of alert generation
  * Suppress state of the alert
  * Acknowledge state of the alert
  * Probable cause for the alert
  * Possible effect due to the alert
  * Recommended corrective actions to follow



``[-node {<nodename>|local}]`` - Node::
Selects the alerts generated for the specified node.

``[-monitor <hm_type>]`` - Monitor::
Selects the alerts with the specified monitor name.
```

```
``[-alert-id <text>]`` - Alert ID::
Selects the alerts with the specified alert ID.

``[-alerting-resource <text>]`` - Alerting Resource::
Selects the alerts with the specified alerting resource name.

``[-subsystem <hm_subsystem>]`` - Subsystem::
Selects the alerts generated on the monitoring subsystem.

``[-indication-time <Date>]`` - Indication Time::
Selects the alerts with the specified indicated time.

``[-perceived-severity <hm_perceived_sev>]`` - Perceived Severity::
Selects the alerts with the perceived severity level.

``[-probable-cause <hm_probable_cause>]`` - Probable Cause::
Selects the alerts that contain the specified probable cause.

``[-probable-cause-description <text>]`` - Description::
Selects the alerts containing the specified probable cause description.

``[-corrective-actions <text>]`` - Corrective Actions::
Selects the alerts with the specified recommended corrective action.

``[-possible-effect <text>]`` - Possible Effect::
Selects the alerts with the specified possible effect.

``[-acknowledge {true|false}]`` - Acknowledge::
Selects the alerts with the specified acknowledgement status.

``[-suppress {true|false}]`` - Suppress::
Selects the alerts with the specified suppressor field status of true or
false.

``[-policy <text>]`` - Policy::
Selects the alerts with the specified policy name.

``[-acknowledger <text>]`` - Acknowledger::
Selects the alerts with the specified acknowledger field.

``[-suppressor <text>]`` - Suppressor::
Selects the alerts with the specified suppressor field.

``[-additional-info <text>,...]`` - Additional Information::
Selects the alerts with the specified additional information.
```

```
``[-alerting-resource-name <text>]`` - Alerting Resource Name::
Selects the alerts with the specified alerting resource name.

``[-tags <hm_alert_type>,...]`` - Additional Alert Tags::
Selects the alerts with the specified keywords.

== Examples

The example below displays information about all the alerts generated in
the cluster:
```

cluster1::> system health alert show
Node: node1
Alert ID: DualPathToDiskShelf_Alert
Resource: Shelf ID 2
Severity: Major
Suppress: false
Acknowledge: false
Tags: quality-of-service, nondisruptive-upgrade
Probable Cause: Disk shelf 2 does not have two paths to controller
node1.
Possible Effect: Access to disk shelf 2 via controller node1 will be
lost with a single hardware component failure (e.g.
cable, HBA, or IOM failure).
Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
2. Connect disk shelf 2 to controller node1 via two paths following the rules in the Universal SAS and ACP Cabling Guide.
3. Reboot the halted controllers.
4. Contact support personnel if the alert persists.

```
The example below displays additional information about a specific alert
generated in the cluster:
```

cluster1::> system health alert show -monitor node-connect -alert-id DualPathToDiskShelf_Alert -instance
Node: node1
Monitor: node-connect
Alert ID: DualPathToDiskShelf_Alert
Alerting Resource: 50:05:0c:c1:02:00:0f:02
Subsystem: SAS-connect
Indication Time: Mon Mar 21 10:26:38 2011
Perceived Severity: Major
Probable Cause: Connection_establishment_error
Description: Disk shelf 2 does not have two paths to controller node1.
Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
2. Connect disk shelf 2 to controller node1 via two paths following the rules in the Universal SAS and ACP Cabling Guide.
3. Reboot the halted controllers.
4. Contact support personnel if the alert persists.
Possible Effect: Access to disk shelf 2 via controller node1 will be lost with a single hardware component failure (e.g. cable, HBA, or IOM failure).

Acknowledge: false
Suppress: false
Policy: DualPathToDiskShelf_Policy
Acknowledger: -
Suppressor: -
Additional Information: Shelf uuid: 50:05:0c:c1:02:00:0f:02
Shelf id: 2
Shelf Name: 4d.shelf2
Number of Paths: 1
Number of Disks: 6
Adapter connected to IOMA:
Adapter connected to IOMB: 4d
Alerting Resource Name: Shelf ID 2
Additional Alert Tags: quality-of-service, nondisruptive-upgrade

```
[[ID8fc755ed4aef476a6c905d41423f6711]]
= system health alert definition show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display system health alert definition

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health alert definition show` command displays information
about the various alerts defined in the system health monitor policy file.
Using `-instance` will display additional details.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Use this parameter to display additional information on each alert
definition.
+
  * Node name
  * Monitor name
  * Subsystem identifier
  * Alert ID
  * Severity of the alert
```

```
  * Probable cause
  * Probable cause description
  * Possible effect due the error state
  * Recommended corrective actions to be followed
  * Any additional information
  * Additional alert tags



``[-node {<nodename>|local}]`` - Node::
Selects the alert definitions for the specified node.

``[-monitor <hm_type>]`` - Monitor::
Selects the alert definitions with the specified monitor name.

``[-alert-id <text>]`` - Class of Alert::
Selects the alert definitions with the specified alert identifier.

``[-perceived-severity <hm_perceived_sev>]`` - Severity of Alert::
Selects the alert definitions with the specified perceived severity.

``[-probable-cause <hm_probable_cause>]`` - Probable Cause::
Selects the alert definitions with the specified probable cause of the
alert.

``[-probable-cause-description <text>]`` - Probable Cause Description::
Selects the alert definitions with the specified probable cause
description.

``[-possible-effect <text>]`` - Possible Effect::
Selects the alert definitions with the specified possible effect.

``[-corrective-actions <text>]`` - Corrective Actions::
Selects the alert definitions with the specified corrective action.

``[-subsystem <hm_subsystem>]`` - Subsystem Name::
Selects the alert definitions with the specified subsystem.

``[-additional-information <text>]`` - Additional Relevant Data::
Selects the alert definitions with the specified additional information.

``[-tags <hm_alert_type>,...]`` - Additional Alert Tags::
Selects the alert definitions with the specified keywords.

== Examples

The example below displays information about all the definitions in the
```

```
alert definition file:
```

cluster1::> system health alert definition show
Node Monitor Subsystem Alert ID
---------- ---------- ---------- ----------
node-01 system-connect SAS-connect DualControllerNonHa_
Alert
Severity: Major
Probable Cause: Configuration_error
Probable Cause Description: Disk shelf $(sschm_shelf_info.id) is connected to
two controllers
($(sschm_shelf_info.connected-nodes)) that are
not an HA pair.
Possible Effect: Access to disk shelf $(sschm_shelf_info.id) may
be lost with a single controller failure.
Corrective Actions: 1. Halt all controllers that are connected to disk shelf $(sschm_shelf_info.id).
2. Connect disk shelf $(sschm_shelf_info.id) to both HA controllers following the rules in the Universal SAS and ACP Cabling Guide.
3. Reboot the halted controllers.
4. Contact support personnel if the alert persists.
Additional Info: -
Tags: quality_of_service, nondisruptive-upgrade

```
The example below displays detailed information about the definitions in
the alert definition file:
```

cluster1::> system health alert definition show -instance
Node: krivC-01
Monitor: system-connect
Class of Alert: DualControllerNonHa_Alert
Severity of Alert: Major
Probable Cause: Configuration_error
Probable Cause Description: Disk shelf $(sschm_shelf_info.id) is connected to two controllers ($(sschm_shelf_info.connected-nodes)) that are not an HA pair.
Possible Effect: Access to disk shelf $(sschm_shelf_info.id) may be lost with a single controller failure.
Corrective Actions: 1. Halt all controllers that are connected to disk shelf $(sschm_shelf_info.id).
2. Connect disk shelf $(sschm_shelf_info.id) to both HA controllers following the rules in the Universal SAS and ACP Cabling Guide.
3. Reboot the halted controllers.
4. Contact support personnel if the alert persists.
Subsystem Name: SAS-connect
Additional Relevant Data: -
Additional Alert Tags: quality_of_service, nondisruptive-upgrade

```
[[ID83873627a74352446d83ad98e93fe0f4]]
= system health autosupport trigger history show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
View system health alert history

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health autosupport trigger history show` command displays all
the alert triggers in the cluster that generated the AutoSupport messages.
The following fields are displayed in the output:

* Node name
* Monitor name
* Subsystem
* Alert identifier
* Alerting resource
* Severity
* If an AutoSupport has been sent due to this alert



== Parameters

{ ``[-fields <fieldname>,...]``::
Use this parameter to display only the fields you specify.

| ``[-instance  ]`` }::
Use this parameter to display additional information about all of the
alerts that were generated.

``[-node {<nodename>|local}]`` - Node::
Use this parameter to display AutoSupport trigger history on the specified
node.

``[-monitor <hm_type>]`` - Monitor::
Use this parameter to display AutoSupport trigger history with the
specified monitor name.

``[-alert-id <text>]`` - Alert ID::
Use this parameter to display the AutoSupport message that was triggered
by the specified alert ID.

``[-alerting-resource <text>]`` - Alerting Resource::
Use this parameter to display the AutoSupport message that was triggered
```

```
by the specified alerting resource.

``[-subsystem <hm_subsystem>]`` - Subsystem::
Use this parameter to display the AutoSupport message that was triggered
by the specified subsystem.

``[-indication-time <Date>]`` - Indication Time::
Use this parameter to display the AutoSupport message that was triggered
at the indicated time.

``[-perceived-severity <hm_perceived_sev>]`` - Perceived Severity::
Use this parameter to display the AutoSupport message that was triggered
by alerts with the specified perceived severity.

``[-autosupport-triggered {true|false}]`` - AutoSupport Triggered::
Use this parameter to display the alerts that generated AutoSupport
messages.

``[-probable-cause <hm_probable_cause>]`` - Probable Cause::
Use this parameter to display the alerts that were generated with the
specified probable cause.

``[-corrective-actions <text>]`` - Corrective Actions::
Use this parameter to display the AutoSupport alerts with the specified
corrective actions.

``[-asup-enable {true|false}]`` - Enable Asup for This Alert::
Use this parameter to enable or disable an AutoSupport message for this
alert.

``[-alert-clear-time <Date>]`` - Alert Clear Time::
Use this parameter to display the alerts that were cleared at a given
time.

== Examples

This example displays information about the AutoSupport trigger history
```

```
cluster1::> system health autosupport trigger history show
Node Monitor Subsystem Alert ID
------------ --------------------- ----------------- ---------------------
node1 node-connect SAS-connect DualPathToDiskShelf_
Alert
Resource: 50:05:0c:c1:02:00:0f:02
Severity: Major
AutoSupport sent: true
```

```
This example displays info about the autosupport trigger history in detail
```

cluster1::> system health autosupport trigger history show -instance
Node: node1
Monitor: node-connect
Alert ID: DualPathToDiskShelf_Alert
Alerting Resource: 50:05:0c:c1:02:00:0f:02
Subsystem: SAS-connect
Indication Time: Thu Mar 17 11:59:09 2011
Perceived Severity: Major
AutoSupport Triggered: true
Probable Cause: Connection_establishment_error
Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
2. Connect disk shelf 2 to controller node1 via two paths following the rules in the Universal SAS and ACP Cabling Guide.
3. Reboot the halted controllers.
4. Contact support personnel if the alert persists.
Enable asup for this alert: true
Alert Clear Time: Wed May 29 16:10:13 2013

```
[[ID4a6b60e33892eef3ff0ef1c155078d31]]
= system health config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display system health configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health config show` command displays the configuration and
status of each health monitor in the cluster. The command shows a health
status for each health monitor. The health status is an aggregation of the
subsystem health for each subsystem that the health monitor monitors. For
example, if a health monitor monitors two subsystems and the health status
of one subsystem is "ok" and the other is "degraded", the health status
for the health monitor is "degraded".

== 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>|local}]`` - Node::
Use this parameter to list the health monitors present on the specified
node.

``[-monitor <hm_type>]`` - Monitor::
Use this parameter to display the health monitors with the specified
monitor name.

``[-subsystem <hm_subsystem>,...]`` - Subsystem::
Selects the health monitors with the specified subsystems.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Health::
Selects the health monitors with the specified health status.

``[-mon-version <text>]`` - Monitor Version::
Selects the health monitors with the specified monitor version.

``[-pol-version <text>]`` - Policy File Version::
Selects the health monitors with the specified health monitor policy
version.

``[-context {Node |Cluster}]`` - Context::
Selects the health monitors with the specified running context.

``[-aggregator <hm_type>]`` - Aggregator::
Selects the health monitors with the specified aggregator.

``[-resources <text>,...]`` - Resource::
Selects the health monitors with the specified resource name.

``[-init-state
{Invalid|Initailizing|Initialized|Starting_Discovery|Starting_Re-
Discovery|Discovery_Done_Partially|Discovery_Done}]`` - Subsystem
Initialization Status::
Selects the health monitors with the specified subsystem initialization
state.

``[-sub-pol-versions <text>]`` - Subordinate Policy Versions::
Selects the health monitors with the specified subordinate policy version.
```

```
== Examples

The example below displays information about health monitor configuration:
```

cluster1::> system health config show
Node Monitor Subsystem Health
---------- --------------------- ---------------- -----------------
node1 node-connect SAS-connect degraded
node1 system-connect SAS-connect degraded
node1 system SAS-connect degraded

```
The example below displays detailed information about health monitor
configuration:
```

cluster1::> system health config show -instance
Node: node1
Monitor: node-connect
Subsystem: SAS-connect
Health: degraded
Monitor Version: 1.0
Policy File Version: 1.0
Context: node_context
Aggregator: system-connect
Resource: SasAdapter, SasDisk, SasShelf
Subsystem Initialization Status: initialized
Subordinate Policy Versions: 1.0 SAS, 1.0 SAS multiple adapters

```
[[ID7ef7f72b898c1ebc6af468e0b091ae6f]]
= system health policy definition modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify system health policy definition

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health policy definition modify` enables or disables health
monitoring policies based on input parameters the user provides.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node on which you want to enable or
disable the policy.

``-monitor <hm_type>`` - Monitor::
Use this parameter to specify the monitor name for which you want to be
enable or disable the policy.

``-policy-id <text>`` - Policy::
Use this parameter to specify the policy identifier that you want to
enable or disable.

``[-enable {true|false}]`` - Policy Status::
Use this parameter with the value "true" to enable the policy. Set the
value to "false" to disable the policy.

``[-asup-enable {true|false}]`` - Enable AutoSupport for This Alert::
Use this parameter to enable or disable an AutoSupport message for this
alert.

== Examples

This example modifies policy state on the cluster:
```

cluster1::> system health policy definition modify -node node1

-policy-id ControllerToShelfIomA_Policy -enable false -monitor *

```
[[IDb635cedaed7f942e7b27a3e7d5942864]]
= system health policy definition show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display system health policy definitions

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health policy definition show` command lists the health
monitor policy definitions as described by the health monitor policy file.
The command displays the following fields:

* Node name
* Monitor name
* Policy name
* Policy rule expression
* Expression for joining two tables
* Policy status
* Alert identifier
* Responsible resource name



== 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>|local}]`` - Node::
Selects policy definitions for the specified node.
```

```
``[-monitor <hm_type>]`` - Monitor::
Selects policy definitions with the specified monitor name.

``[-policy-id <text>]`` - Policy::
Selects policy definitions with the specified policy identifier.

``[-rule-expression <ArithExpr>]`` - Rule Expression::
Selects policy definitions with the specified rule of expression.

``[-where <ArithExpr>]`` - Variable Equivalence::
Selects rules that match the provided expression. This expression is part
of the alert definition. It is shown for reference only and cannot be
changed.

``[-enable {true|false}]`` - Policy Status::
Use this parameter with the value set to "true" to select policy
definitions that are enabled. Set the value to "false" to select policy
definitions that are disabled.

``[-alert-id <text>]`` - Alert ID::
Selects all policy definitions of the specified alert identifier.

``[-responsible-resource-info <text>]`` - Table and ID of Resource at
Fault::
Selects all policy definitions with the specified responsible resource.

``[-asup-enable {true|false}]`` - Enable AutoSupport for This Alert::
Selects policy definitions for which AutoSupport messages are either
enabled or disabled.

== Examples

The example below displays information about all the policy definitions
present in the cluster:
```

cluster1::> system health policy definition show
Node Monitor Policy
--- --- ---
node1 node-connect ControllerToShelfIomA_Policy
Policy Rule Expression: nschm_shelf_info.num-paths == 2
nschm_shelf_info.iomb-adapter == NULL
Where: -
Enable: true
Alert ID: ControllerToShelfIomA_Alert
Number of Alerts: -
Responsible Resource: nschm_shelf_info.name

```
The example below displays detailed information about all the policy
definitions present in the cluster:
```

cluster1::> system health policy definition show -instance
Node: node1
Monitor: node-connect
Policy: ControllerToShelfIomA_Policy
Rule Expression: nschm_shelf_info.num-paths == 2 nschm_shelf_info.iomb-adapter == NULL
Variable Equivalence: -
Policy Status: true
Alert ID: ControllerToShelfIomA_Alert
Table and ID of Resource at Fault: nschm_shelf_info.name

```
[[ID67decf6dd3159151cf138fe9257d4c16]]
= system health status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display system health monitoring status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health status show` command displays the health monitor
status. The possible states are:

* ok
* ok-with-suppressed
* degraded
* unreachable



== Examples

This example displays information about health monitoring status:
```

cluster1::> system health status show
Status
---
degraded

```
[[IDed199298d23063c5998f2f9e71da52b7]]
= system health subsystem show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the health of subsystems

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system health subsystem show` command displays the health status of
each subsystem for which health monitoring is available. This command
aggregates subsystem health status from each node in the cluster. A
subsystem's health status changes to "degraded" when a health monitor
raises an alert. You can use the xref:{relative_path}system-health-alert-
show.html[system health alert show] command to display information about
generated alerts.

== 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.

``[-subsystem <hm_subsystem>]`` - Subsystem::
Selects the specified subsystem.

``[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Health::
Selects subsystems that have the specified health status.

``[-init-state
{Invalid|Initailizing|Initialized|Starting_Discovery|Starting_Re-
Discovery|Discovery_Done_Partially|Discovery_Done}]`` - Initialization
State::
Selects subsystems that have the specified initialization state.
```

```
``[-outstanding-alert-count <integer>]`` - Number of Outstanding Alerts::
Selects subsystems that have the specified number of outstanding alerts.

``[-suppressed-alert-count <integer>]`` - Number of Suppressed Alerts::
Selects subsystems that have the specified number of suppressed alerts.

``[-node {<nodename>|local}]`` - Node::
Selects subsystems for the specified node.

``[-refresh-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>,...]`` - Subsystem
Refresh Interval::
The refresh interval is in minutes. A value of zero disables the sub-
system refresh until a reboot or restart of the subsystem process.

== Examples

The example below displays the health status of each subsystem:
```

cluster1::> system health subsystem show
Subsystem Health
---------- ----------
SAS-connect degraded
Switch-Health OK
CIFS-NDO OK

```
The example below displays detailed information about the health status of
each subsystem:
```

cluster1::> system health subsystem show -instance

```
                              Subsystem: SAS-connect
                                 Health: degraded
                   Initialization State: initialized
           Number of Outstanding Alerts: 0
            Number of Suppressed Alerts: 0
                                   Node: node1,node2
             Subsystem Refresh Interval: 30m, 30m
Subsystem: Switch-Health
                                 Health: ok
                   Initialization State: initialized
           Number of Outstanding Alerts: 0
            Number of Suppressed Alerts: 0
                                   Node: node1
             Subsystem Refresh Interval: 5m
Subsystem: CIFS-NDO
                                 Health: OK
                   Initialization State: initialized
           Number of Outstanding Alerts: 0
            Number of Suppressed Alerts: 0
                                   Node: node1
             Subsystem Refresh Interval: 5m
```

```
== Related Links

* xref:{relative_path}system-health-alert-show.html[system health alert
show]




:leveloffset: -1


= system license commands

:leveloffset: +1


[[ID4c0c99a7d4dad978de302f68db5e582e]]
= system license add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Add one or more licenses

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command adds a license to a cluster. To add a license you must
specify a valid license key, which you can obtain from your sales
representative.

== Parameters

``-license-code <License Code V2>,...`` - License Code V2::
This parameter specifies the key of the license that is to be added to the
cluster. The parameter accepts a list of 28 digit upper-case alphanumeric
character keys.

| ``[-use-license-file {true|false}]`` - Use License File [.advaced-
priv]#(privilege: advanced)# }::
If this parameter is set to true, licenses from the local node
/mroot/etc/lic_file license file will be installed if the file exists at
this location.

== Examples

The following example adds a list of licenses with the keys
AAAAAAAAAAAAAAAAAAAAAAAAAAAA and BBBBBBBBBBBBBBBBBBBBBBBBBBBB to the
cluster
```

cluster1::> system license add -license-code AAAAAAAAAAAAAAAAAAAAAAAAAAAA, BBBBBBBBBBBBBBBBBBBBBBBBBBBB

```
The following example installs the licenses from the local node
"/mroot/etc/lic_file"
```

cluster1::> system license add -use-license-file true

```
[[ID31aff617141ae81677b82fd7a9516051]]
= system license clean-up
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove unnecessary licenses

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command manages licenses in the cluster that have no effect, and so
can be removed. Licenses that have expired or are not affiliated with any
controller in the cluster are deleted by this command. Licenses that
cannot be deleted are displayed with reasons for the non-deletion.

== Parameters

``[-unused <true>]`` - Remove unused licenses::
If you use this parameter, the command removes licenses in the cluster
that are not affiliated with any controller in the cluster.

``[-expired <true>]`` - Remove expired licenses::
If you use this parameter, the command removes licenses in the cluster
that have expired.

``[-n, -simulate <true>]`` - Simulate Only::
If you use this parameter, the command will not remove the licenses.
Instead it will display the licenses that will be removed if this
parameter was not provided.

== Examples

The following example simulates and displays the licenses that can be
cleaned up:
```

cluster-1::> system license clean-up -expired -unused

The following licenses were safely deleted:

Serial number: 1-80-000011
Owner: cdancluster-1

```
Package Reason
------------------------- -----------------------------------------------------
CIFS License has expired

Serial number: 4067154888
Owner: none
Package Reason
------------------------- -----------------------------------------------------
Cloud License has expired

Serial number: 1-81-0000000000000000004067154999
Owner: none
Package Reason
------------------------- -----------------------------------------------------
iSCSI License unused by any node in the cluster
The following licenses are either expired or unused but cannot be safely deleted:

Serial number: 4067154778
Owner: node1
Package Reason
------------------------- -----------------------------------------------------
Cloud Feature would be impaired upon removal

Serial number: 4067154779
Owner: node2
Package Reason
------------------------- -----------------------------------------------------
Cloud System generated license
```

```
[[ID43c35ac923c27052d078288bffe75c8b]]
= system license delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a license

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command deletes a license from a cluster.

== Parameters

``-serial-number <text>`` - Serial Number::
This parameter specifies the serial number of the license that is to be
deleted from the cluster. If this parameter is not provided, the default
value is the serial number of the cluster.

``-package <Licensable Package>`` - Package::
This parameter specifies the name of the package that is to be deleted
from the cluster.

== Examples

The following example deletes a license named CIFS and serial number 1-81-
0000000000000000000123456 from the cluster:
```

cluster1::> system license delete -serial-number 1-81-0000000000000000000123456 -package CIFS

```
The following example deletes from the cluster all of the licenses under
the installed-license Core Bundle for serial number 123456789:
```

cluster1::> system license delete { -serial-number 123456789 -installed-license "Core Bundle" }

```
[[IDdbdfc5ca155b95b82d5cd767430ee709]]
= system license show-aggregates
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status of aggregates leases and license used.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the status of all ONTAP aggregates.

== 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>|local}]`` - Node::
If you use this parameter, the command displays information only about
aggregates that match the given node.

``[-aggr-name <text>]`` - Aggregate Name::
If you use this parameter, the command displays information only about
aggregate that match the given aggregate.

``[-aggr-size {<integer>[KB|MB|GB|TB|PB]}]`` - Aggregate Size::
If you use this parameter, the command displays information only about
aggregates that match the given physical size of an aggregate.

``[-licensed-size {<integer>[KB|MB|GB|TB|PB]}]`` - Licensed Size::
If you use this parameter, the command displays information only about
aggregates that match the given licensed-size.

``[-expiration <MM/DD/YYYY HH:MM:SS>]`` - Lease Expiration::
If you use this parameter, the command displays information only about
aggregates that match the given lease expiration.

``[-status <AggrLicStatus>]`` - Aggregate Status::
If you use this parameter, the command displays information only about
aggregates that match the given status.
```

```
``[-compliant {true|false}]`` - Is Aggregate Compliant::
If you use this parameter, the command displays information only about
aggregates that match the given state of compliance.

``[-aggr-uuid <UUID>]`` - Aggregate UUID::
If you use this parameter, the command displays information only about
aggregate that match the given aggregate uuid.

== Examples

The following example displays the license status of the cluster:
```

cluster1::> system license show-aggregates
Licensed Physical
Node Aggregate Size Size Lease Expiration Status
--------- ------------------ -------- -------- ------------------ ------------
node1
root1 0B 2GB - lease-not-required
root2 (mirror) 0B 2GB - lease-not-required
aggr1 20GB 20GB 6/21/2018 18:10:00 lease-up-to-date
aggr2 (mirror) 10GB 10GB 6/21/2018 20:00:00 lease-up-to-date
node2
root1 (mirror) 0B 2GB - lease-not-required
root2 0B 2GB - lease-not-required
aggr1 (mirror) 20GB 20GB 6/21/2018 18:10:00 lease-up-to-date
aggr2 10GB 10GB 6/21/2018 20:00:00 lease-up-to-date
node3
root3 0B 2GB - lease-not-required
root4 (mirror) 0B 2GB - lease-not-required
aggr3 15GB 0B 6/21/2018 20:00:00 aggregate-deleted
aggr4 (mirror) 15GB 15GB 6/21/2018 12:00:00 lease-expired
aggr5 (mirror) 15GB 15GB 6/21/2018 21:00:00 lease-up-to-date
aggr6 15GB 15GB 6/21/2018 21:00:00 plex-deleted
aggr7 15GB 14GB 6/21/2018 21:00:00 aggregate-license-size-decreased
aggr8 (mirror) 0B 14GB - lease-missing
node4
root3 (mirror) 0B 2GB - lease-not-required
root4 0B 2GB - lease-not-required
aggr3 (mirror) 15GB 0B 6/21/2018 20:00:00 aggregate-deleted
aggr4 15GB 15GB 6/21/2018 12:00:00 lease-expired
aggr5 15GB 15GB 6/21/2018 21:00:00 lease-up-to-date
aggr6 (mirror) 15GB 0B 6/21/2018 21:00:00 plex-deleted
aggr7 (mirror) 15GB 14GB 6/21/2018 21:00:00 aggregate-license-size-decreased
aggr8 0B 14GB - lease-missing

```
[[ID8fea1094f5fee1f0ee0e11dcd3b51f66]]
= system license show-serial-numbers
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display History of Serial Numbers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the history of changes to the support and node
serial numbers. The node serial number and the support serial number of an
ONTAP system are generally the same and do not change over time. However,
when capacity pools licensing is used, the support serial number is that
of the capacity pool license serial number and the node serial number is
generated by the license manager. Also, when a cluster is upgraded or
converted from capacity tiers licensing to capacity pools licensing, its
support serial numbers as well as its node serial numbers change.

== 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.

``[-seqnum <Sequence Number>]`` - Sequence number::
Event sequence number

``[-node {<nodename>|local}]`` - Node::
Selects the node names that match this parameter value.

``[-date <MM/DD/YYYY HH:MM:SS>]`` - Date::
Selects the dates of serial number changes that match this parameter
value.

``[-reason <text>]`` - Reason for change::
Reasons for serial number changes.

``[-support-serial-number <text>]`` - Support Serial Number::
Selects the support serial numbers that match this parameter value.
```

```
``[-node-serial-number <text>]`` - Node Serial Number::
Selects the node serial number that match this parameter value.

== Examples

The following example displays the serial number change history of a four
node capacity pools cluster. Its two HA pairs were originally assigned to
capacity pools 390000101 and 390000102, and then both were reassigned to
another capacity pool 390000103:
```

cluster1::> system license show-serial-numbers

Change Date: 06/06/2019
Reason: Reassignments of capacity pools
Node Support Serial Node Serial
--- --- ---
node1 390000103 99939000010100000001
node2 390000103 99939000010100000002
node3 390000103 99939000010200000003
node4 390000103 99939000010200000004

Change Date: 03/01/2019
Reason: Initial installation
Node Support Serial Node Serial
--- --- ---
node1 390000101 99939000010100000001
node2 390000101 99939000010100000002
node3 390000102 99939000010200000003
node4 390000102 99939000010200000004

```
The following example displays the history of a four node cluster
converted from capacity tiers licensing to capacity pools licensing:
```

cluster1::> system license show-serial-numbers

Change Date: 06/06/2019
Reason: Conversions from capacity tiers to pools
Node Support Serial Node Serial
--- --- ---
node1 390000103 99939000010300000011
node2 390000103 99939000010300000012
node3 390000103 99939000010300000013
node4 390000103 99939000010300000014

Change Date: 03/01/2019
Reason: Initial installation
Node Support Serial Node Serial
--- --- ---
node1 310000101 310000101
node2 310000102 310000102

node3 310000103 310000103
node4 310000104 310000104

```
The following example displays the history of an evaluation cluster that
was upgraded to capacity pools licensing:
```

cluster1::> system license show-serial-numbers

Change Date: 06/06/2019
Reason: Conversions from capacity tiers evaluation to pools
Node Support Serial Node Serial
--- --- ---
node1 390000103 99939000010300000011
node2 390000103 99939000010300000012
node3 390000103 99939000010300000013
node4 390000103 99939000010300000014

Change Date: 03/01/2019
Reason: Initial installation
Node Support Serial Node Serial
--- --- ---
node1 evaluation 99887766554433221101
node2 evaluation 99887766554433221102
node3 evaluation 99887766554433221103
node4 evaluation 99887766554433221104

```
[[ID2f1d8c3ca8410b7a433a586b2aaf91b5]]
= system license show-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display license status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the status of all ONTAP licenses.

== 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.

``[-status {not-compliant|eval|partially-installed|valid|not-
installed|not-applicable|not-known}]`` - Current State::
If you use this parameter, the command displays information only about
licenses that match the given status.

``[-license <Licensable Package>]`` - License::
If you use this parameter, the command displays information only about
licenses that match the given license.

``[-scope {site|cluster|node|pool}]`` - License Scope::
If you use this parameter, the command displays information only about
licenses that match the given scope.

``[-detailed-status <text>,...]`` - Detailed Status::
If you use this parameter, the command displays information only about
licenses that match the given detailed-status.

== Examples

The following example displays the license status of the cluster:
```

cluster1::> system license show-status
Status License Scope Detailed Status
--------- ------------------- --------- ----------------------
partially-installed
CIFS node License missing on: Node2-Cluster1.
SnapRestore node License missing on: Node2-Cluster1.
valid
FCP node -
FabricPool cluster The system is using 1TB, and can use up to 25TB.
not-installed
NFS - -
iSCSI - -
SnapMirror - -
FlexClone - -
SnapVault - -
SnapLock - -
SnapManagerSuite - -
SnapProtectApps - -
V_StorageAttach - -
Insight_Balance - -
OCShift - -
TPM - -

VE - -
DP_Optimized - -
not-applicable
Cloud - -
Select - -
20 entries were displayed.

```
[[IDc7b38242f134b820b5446411e8932347]]
= system license show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display licenses

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The ` system license show ` command displays the information about
licenses in the system.

== 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.

``[-serial-number <text>]`` - Serial Number::
If you use this parameter, the command displays information only about the
licenses that matches the serial number you specify.

``[-package <Licensable Package>]`` - Package::
If you use this parameter, the command displays information only about the
specified package.

``[-owner <text>]`` - Owner::
If you use this parameter, the command displays information only about the
packages that matches the owner name you specify.
```

```
``[-expiration <MM/DD/YYYY HH:MM:SS>]`` - Expiration::
If you use this parameter, the command displays information only about the
licenses that have the expiration date you specify.

``[-description <text>]`` - Description::
If you use this parameter, the command displays information only about the
licenses that matches the description you specify.

``[-type {license|site|demo|subscr|capacity|capacity-per-term|enabled}]``
- Type::
If you use this parameter, the command displays information only about the
licenses that have the license type you specify.

``[-customer-id <text>]`` - Customer ID::
If you use this parameter, the command displays information only about the
licenses that have the customer-id you specify.

``[-installed-license <text>]`` - Installed License Name::
If you use this parameter, the command displays information only about the
licenses that match the installed license you specify.

``[-host-id <text>]`` - Host Id::
If you use this parameter, the command displays information only about the
license that have the host id you specify.

``[-capacity {<integer>[KB|MB|GB|TB|PB]}]`` - License Capacity::
If you use this parameter, the command displays information only about the
licenses that match the capacity you specify.

== Examples

The following example displays information about all licensed packages in
the cluster:
```

cluster1::> system license show

Serial Number: 1-81-0000000000000001122334455
Owner: node2
Installed License: Legacy Key
Capacity: -
Package Type Description Expiration
---------- -------- ---------- ----------
NFS license NFS License -
CIFS license CIFS License -
iSCSI license iSCSI License -
SnapRestore license SnapRestore License -
FlexClone license FlexClone License -

S3 license S3 License -

Serial Number: 123456789
Owner: node1
Installed License: Core Bundle
Capacity: 10TB
Package Type Description Expiration
---------- -------- ---------- ----------
NFS capacity NFS License -
CIFS capacity CIFS License -
iSCSI capacity iSCSI License -
SnapRestore capacity SnapRestore License -
FlexClone capacity FlexClone License -
S3 capacity S3 License -
12 entries were displayed.

```
[[IDab62a29c42829f55662beb9c99c11a7f]]
= system license update-leases
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Begin lease reconciliation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The ` system license update-leases` command attempts to update (that is,
renew) any capacity pool leases that have expired.

== Parameters

``[-node {<nodename>|local}]`` - Nodes to Attempt Renewal::
This optional parameter directs the system to update leases for only the
specified nodes.

``[-force {true|false}]`` - Force Renewal of Valid Leases::
This optional parameter, if set with a value of "true", directs the system
to update all leases for a node, not just those that have expired.

== Examples

The following example updates all leases on a node:
```

cluster1::*> system license update-leases -node node1 -force true
Number of Leases Updated: 3
Number of Leases Not Updated: 0 (error), 0 (already up-to-date)

```
[[ID4f51b4a0631b2e2ffa4c31fd1c933771]]
= system license capacity show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Show license capacity status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future release of
ONTAP. Use the "link:system-license-show-status.html[system license show-
status] " command.

The `system license capacity show` command displays the information about
the licenses in the system that are specifically related to storage
capacity limits.

== 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.

``[-serial-number <Node Serial Number>]`` - Serial Number::
If you use this parameter, the command displays information only about the
capacity-related licenses that matches the serial number you specify.

``[-package <Licensable Package>]`` - Package::
If you use this parameter, the command displays information only about the
package you specify.
```

```
``[-owner <text>]`` - Owner::
If you use this parameter, the command displays information only about the
capacity-related licenses that have the owner you specify.

``[-max-capacity {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Capacity::
If you use this parameter, the command displays information only about the
capacity-related licenses that have the maximum amount of attached storage
capacity you specify.

``[-current-capacity {<integer>[KB|MB|GB|TB|PB]}]`` - Current Capacity::
If you use this parameter, the command displays information only about the
capacity-related licenses that apply to the node with the current attached
capacity you specify.

``[-expiration <MM/DD/YYYY HH:MM:SS>]`` - Expiration Date::
If you use this parameter, the command displays information only about the
capacity-related licenses that have the expiration date you specify.

``[-reported-state {evaluation|warning|missing|enforcement|installed}]`` -
Reported State::
If you use this parameter, the command displays information only about the
capacity-related licenses that have the reported state you specify.

``[-node {<nodename>|local}]`` - Node Name::
If you use this parameter, the command displays information only about the
capacity-related licenses that apply to the node you specify.

== Examples

The following example displays information about all capacity-related
licensed packages in the cluster, for a hypothetical cluster of four
nodes:

Note that for some nodes below, the maximum capacity is displayed as "-"
(meaning "unlimited"). This happens when there is no capacity license for
the node - the node is operating with a limited-time temporary capacity
license.
```

cluster1::> system license capacity show

Node: node1
Serial Number: 1-81-0000000000001234567890123456
Max Current
Package Capacity Capacity Expiration
---------- -------- -------- ----------
Select 2TB 15.81GB 4/11/2016 00:00:00
Node: node2

Serial Number: 1-81-0000000000000000000123456788
Max Current
Package Capacity Capacity Expiration
---------- -------- -------- ----------
Select - 10.40TB 4/11/2016 00:00:00
Node: node3
Serial Number: 1-81-0000000000000000000123456789
Max Current
Package Capacity Capacity Expiration
---------- -------- -------- ----------
Select - 10.40TB 4/11/2016 00:00:00
Node: node4
Serial Number: 1-81-0000000000000123456789012345 6
Max Current
Package Capacity Capacity Expiration
---------- -------- -------- ----------
Select 2TB 15.81GB 4/11/2016 00:00:00

```
== Related Links

* xref:{relative_path}system-license-show-status.html[system license show-
status]




[[IDc4d12d337b80ef1585fd88f595d1da7c]]
= system license entitlement-risk show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Cluster License Entitlement Risk

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays information about license entitlement risk of the
cluster for each license package. The command displays license package
name, entitlement risk, corrective action to reduce the entitlement risk
for each package, and the names and serial numbers for the nodes that do
not have a node-locked license for a given package. If command is used
with the "-detail" parameter, the output displays the names and serial
numbers for all nodes in the cluster instead of only the nodes missing a
```

```
node-locked license. It also displays whether each node has a license and
if the features corresponding to the package are used in the past week.



License entitlement risk does not apply to base license. If a node has a
site or a valid demo license for the given package, the entitlement risk
will be shown as "medium" and the nodes missing a node-locked license will
be displayed. The corrective action, if the cluster has a site license for
the given package is, "Verify all controllers are entitled". If the
entitlement risk is high, the corrective action is "Acquire node-locked
license". For the low entitlement risk and if the cluster is unlicensed
for a given package, the corrective action is "None". If the license
entitlement risk cannot be computed because of infrastructure issues, the
entitlement risk is shown as "unknown" and the corrective action is
displayed as "Verify system health".For more information regarding license
entitlement risk, see +
http://mysupport.netapp.com/licensing/ontapentitlementriskstatus

== Parameters

{ ``[-fields <fieldname>,...]``::
With this parameter, you can specify which fields should be displayed by
the command. License package names and node serial numbers are always
displayed.

| ``[-detail  ]``::
If you use this parameter, the command displays the license package name,
entitlement risk, corrective action, all nodes' names, their serial
numbers, whether a node-locked license is present and whether a given
license package has been in use in the past week for each node in the
cluster.

| ``[-instance  ]`` }::
If this parameter is used, the command displays values for all fields for
each license package and each node in the cluster individually.

``[-package <Licensable Package>]`` - Package Name::
If you use this parameter, the command displays information only for the
specified license package.

``[-serial-number <text>]`` - Node Serial Number::
If you use this parameter, the command displays information only for the
node with the specified serial number. The displayed entitlement risk and
corrective action apply to the entire cluster.

``[-node-name <text>]`` - Node Name::
```

```
If you use this parameter, the command displays information only for the
node with the specified name. The displayed entitlement risk and
corrective action apply to the entire cluster.

``[-risk {high|medium|low|unlicensed|unknown}]`` - Entitlement Risk::
If you use this parameter, the command displays information only for the
license packages that have the specified license entitlement risk.

``[-action <text>]`` - Corrective Action::
If you use this parameter, the command displays information only for the
license packages which need the specified corrective action to reduce
entitlement risk.

``[-is-licensed {true|false}]`` - Is Node-Locked License Present::
If you use this parameter, the command displays information only for the
license packages for which at least one node in the cluster has a node-
locked license. It also displays the nodes in the cluster which do not
have a node-locked license.

``[-in-use {true|false}]`` - Usage Status::
If you use this parameter, the command displays information only for the
license packages with corresponding features in use.

``[-missing-serial-numbers <text>,...]`` - Serial Numbers Missing a Node-
Locked License::
If you use this parameter, the command displays the packages for which the
node with the specified serial number does not have a node-locked license.

``[-missing-node-names <text>,...]`` - Node Names Missing a Node-Locked
License::
If you use this parameter, the command displays all the packages for which
the node with the specified name does not have a node-locked license.

``[-action-code {acquire-license|adjust-capacity|verify-
entitlement|verify-system-health|none}]`` - Corrective Action Code::
If you use this parameter, the command displays information only for the
license packages which need specified corrective action code to reduce
entitlement risk. This parameter is same as the parameter "action".

== Examples

The following example displays the information for license package NFS.
NFS is unlicensed in the cluster and no action is necessary to reduce the
entitlement risk. The nodes, cluster1-01 and cluster-02, are missing a
node-locked license. The serial numbers for both nodes are also displayed.
```

cluster1::> system license entitlement-risk show
Package Entitlement Risk Corrective Action
--- --- ---
NFS unlicensed None
Nodes Without a Node-Locked License
---
cluster1-01 1-81-0000000000000004073806282
cluster1-02 1-81-0000000000000004073806283

```
The following example displays the information for license package CIFS.
The cluster has high entitlement risk for CIFS. The command displays
serial numbers for all nodes in the cluster. Both nodes are missing a
node-locked CIFS license. Node with serial number 1-81-
0000000000000004073806282 has used CIFS feature in the past week, and the
node with serial number 1-81-0000000000000004073806283 has not used this
feature in the past week.
```

cluster1::> system license entitlement-risk show -detail
Package Entitlement Risk Corrective Action
--- --- ---
CIFS high Acquire a node-locked license
Serial Numbers Licensed Usage
--- --- ---
1-81-0000000000000004073806282 false true
1-81-0000000000000004073806283 false false

```
[[ID8f1297538346218f1e4105210db55bea]]
= system license license-manager check
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display license manager status

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` system license license-manager check ` checks the connectivity
status of a node to the License Manager that the node was configured to
use. The status of a node might indicate that the License Manager is
inaccessible. If so, the status message contains additional text in
parentheses. The text options and descriptions are as follows:
```

```
* license_expired : The License Manager has a license, but it is expired.
* network_error : The node is unable to establish basic network
connectivity.
* no_valid_license : The License Manager does not have a valid capacity
pool license.

All other values indicate an internal error.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter directs the system to display results for the License
Manager configured for the specified node.

``[-status <text>]`` - Status [.advaced-priv]#(privilege: advanced)#::
This parameter directs the system to display results for the given status
message.

== Examples

The following examples check the status of the configured License Manager,
before and after its license has expired:
```

cluster1::*> system license license-manager check -node node1
Node: node1
LM status: License Manager (1.2.3.4:5678) is accessible.

cluster1::*> system license license-manager check

Node Status
--- ---
node1 License Manager (1.2.3.4:5678) is accessible.
node2 License Manager (1.2.3.4:5678) is accessible.
2 entries were displayed.
cluster1::*> system license license-manager check -node node1
Node: node1
LM status: License Manager (1.2.3.4:5678) is inaccessible (license_expired).

cluster1::*> system license license-manager check

Node Status
--- ---
node1 License Manager (1.2.3.4:5678) is inaccessible (license_expired).
node2 License Manager (1.2.3.4:5678) is inaccessible (license_expired).
2 entries were displayed.

```
[[ID4f2a68c94a3aa648797d9f45c018b8d3]]
= system license license-manager modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]


*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` system license license-manager modify ` command modifies the
configuration information for the License Manager the system is using.

== Parameters

``[-host <text>]`` - License Manager Host [.advaced-priv]#(privilege:
advanced)#::
Sets the specified host, which can either be a fully qualified domain name
(FQDN) or an IP address.

== Examples

The following example modifies information about the License Manager
configuration:
```

cluster1::*> system license license-manager modify -host myhost.mycompany.com

```
[[IDf259b5d28cea3492d417406782f5f516]]
= system license license-manager show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display license manager information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` system license license-manager show ` command displays the
information about the current License Manager configuration.

== Examples

The following example displays information about current License Manager
configuration:
```

cluster1::*> system license license-manager show

License Manager Host: 1.2.3.4

```
[[ID89f1146ff4f26038c8fdb53cc36b85de]]
= system license status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Display license status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future release of
ONTAP. Use the "link:system-license-show-status.html[system license show-
```

```
status] " command.

This command displays the list of licensable packages in the system and
their current licensing status.

== 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.

``[-package <Licensable Package>]`` - Package Name::
If you use this parameter, the command displays information only about the
specified package.

``[-method {none|license|site|demo|subscr|capacity|enabled}]`` - Licensed
Method::
If you use this parameter, the command displays information only about the
packages with the specified licensed method.

``[-expiration <MM/DD/YYYY HH:MM:SS>]`` - Expiration Date::
If you use this parameter, the command displays information only about the
licenses that have the expiration date you specify.

``[-description <text>]`` - Description::
If you use this parameter, the command displays information only about the
licenses that match the description you specify.

``[-status-details <text>]`` - Additional Information About Status::
This option displays additional information regarding the cluster-level
license status for license methods.

== Examples

The following example displays the license status of the cluster:
```

cluster1::> system license status show
Package Licensed Method Expiration Status Details
---------- ---------- ---------- ----------
Base site - -
NFS site - -
CIFS demo 12/7/2015 00:00:00 Demo expires on given date

iSCSI none - -
FCP none - -
SnapRestore none - -
SnapMirror none - -
FlexClone none - -
SnapVault none - -
SnapLock none - -
SnapManagerSuite none - -
SnapProtectApps none - -
V_StorageAttach none - -
SnapLock_Enterprise
none - -
Insight_Balance none - -
OCShift none - -
Cloud subscr 12/15/2015 00:00:00 Subscription expires on given date
17 entries were displayed.

```
== Related Links

* xref:{relative_path}system-license-show-status.html[system license show-
status]




:leveloffset: -1


= system limits commands

:leveloffset: +1


[[ID0406a9c95b7e94362232daf0f8d1c993]]
= system limits show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays the Maximum volume limit

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description
```

```
This command displays the Maximum volumes allowed in a node and maximum
allowed size of a volume.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - Node's filer ID [.advaced-priv]#(privilege:
advanced)#::
This parameter indicates the node name.

``[-max-aggr-size <integer>]`` - Maximum Aggregate Size in TB [.advaced-
priv]#(privilege: advanced)#::
This parameter indicates the Maximum aggregate size in TB

``[-max-node-vvols <integer>]`` - Maximum Number of Volumes [.advaced-
priv]#(privilege: advanced)#::
This parameter indicates the Maximum Number of Volumes

``[-max-hya-ssd-limit <integer>]`` - Maximum SSD on a Flash Pool Aggregate
[.advaced-priv]#(privilege: advanced)#::
This parameter indicates the Maximum hya ssd limit

``[-is-200tb-vvol-size-allowed {true|false}]`` - Is system allowed to
create vvols upto 200TB [.advaced-priv]#(privilege: advanced)#::
This parameter indicates the whether the volume of size 200TB is allowed

``[-max-vol-size-old <integer>]`` - Old Maximum Volume Size in TB
[.advaced-priv]#(privilege: advanced)#::
This parameter indicates Maximum Volume Size in TB prior to 9.12.1

``[-max-vol-size-published <integer>]`` - User Visible Maximum Volume Size
in TB [.advaced-priv]#(privilege: advanced)#::
This parameter indicates Maximum Volume Size in TB published in customer
documentation.

== Examples

Example shows the output of system limits show
```

cluster1::> system limits show
Node Max_volume_size Max_volume_count
---------- ---------- ----------
node1 300 500

```
:leveloffset: -1


= system node commands

:leveloffset: +1


[[ID2b503ab1cab55bc530ecce9f76b71d87]]
= system node halt
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Shut down a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node halt` command stops all activity on a node. You may
supply a reason for the shutdown, which will be stored in the audit log.
You may also keep partner nodes from performing storage takeover during
the shutdown.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this mandatory parameter to specify the node that you want to shut
down. The value `local` specifies the current node.

``[-reason <text>]`` - Reason for Shutdown::
Use this parameter to enter a brief note to indicate the reason for the
restart, which will be stored in the audit log. Providing this information
assists support personnel with troubleshooting efforts.

``[-f, -inhibit-takeover <true>]`` - Disallow Storage Takeover by
Partner::
This parameter optionally forces the shutdown and prevents storage
```

```
failover. LIFs will migrate prior to shutdown even when inhibit-takeover
is set to `true` . To prevent LIF migration, skip-lif-migration-before-
shutdown should be set to `true` . In a two-node MetroCluster
configuration, this parameter prevents automatic unplanned switchover.
+
NOTE: If -inhibit-takeover is set to true, the default behavior of the
xref:{relative_path}storage-failover-show.html[storage failover show]`
-fields onreboot` command is ignored.
+

+
If you enter this command without using this parameter, its effective
value is false and storage takeover is allowed. If you enter this
parameter without a value, it is automatically set to true and storage
takeover is disabled during reboot.

``[-d, -dump <true>]`` - Create a Core Dump::
If this parameter is set to true, it forces a dump of the kernel core when
halting the node.

``[-skip-lif-migration-before-shutdown <true>]`` - Skip Migrating LIFs
Away from Node Prior to Shutdown::
If this parameter is specified, LIF migration prior to the shutdown will
be skipped. However if LIFs on this node are configured for failover,
those LIFs may still failover after the shutdown has occurred. The default
is to migrate LIFs prior to the shutdown. In the default case, the command
attempts to synchronously migrate data and cluster management LIFs away
from the node prior to shutdown. If the migration fails or times out, the
shutdown will be aborted.

``[-ignore-quorum-warnings <true>]`` - Skip Quorum Check Before Shutdown::
If this parameter is specified, quorum checks will be skipped prior to the
shutdown. The operation will continue even if there is a possible data
outage due to a quorum issue.

``[-ignore-strict-sync-warnings <true>]`` - Skip SnapMirror Synchronous
Strict Sync Check Before Reboot::
If this parameter is specified, the check for volumes that are in
SnapMirror Synchronous relationships with policy of type strict-sync-
mirror will be skipped. The operation will continue even if there is a
possible data outage due to not being able to fully sync data.

``[-power-off {true|false}]`` - Power off the node after shutdown::
If this parameter is specified, the node will be powered off at the end of
the halt operation. If set to false, the OS will reboot and then stop at
the LOADER prompt. If set to true, when the node is powered back on, it
```

```
will stop at the LOADER prompt.

== Examples

The following example shuts down the node named cluster1 for hardware
maintenance:
```

cluster1::> system halt -node cluster1 -reason 'hardware maintenance'

```
== Related Links

* xref:{relative_path}storage-failover-show.html[storage failover show]




[[ID3a9eb87b070c48cc31fa0fd1db28df7d]]
= system node migrate-root
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start the root aggregate migration on a node

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node migrate-root` command migrates the root aggregate of a
node to a different set of disks. You need to specify the node name and
the list of disks on which the new root aggregate will be created. The
command starts a job that backs up the node configuration, creates a new
aggregate, set it as new root aggregate, restores the node configuration
and restores the names of original aggregate and volume. The job might
take as long as a few hours depending on time it takes for zeroing the
disks, rebooting the node and restoring the node configuration.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Specifies the node that owns the root aggregate that you want to migrate.
The value ``_local_`` specifies the current node.
```

```
{ ``-disklist <disk path name>,...`` - List of Disks for New Root
Aggregate [.advaced-priv]#(privilege: advanced)#::
Specifies the list of disks on which the new root aggregate will be
created. All disks must be spares and owned by the same node. Minimum
number of disks required is dependent on the RAID type.

``-raid-type {raid_tec|raid_dp|raid4|raid_ep}`` - RAID Type for the New
Root Aggregate [.advaced-priv]#(privilege: advanced)#::
Specifies the RAID type of the root aggregate. The default value is
``_raid-dp_`` .

| ``-resume <true>`` - Resume a Failed Migrate Operation [.advaced-
priv]#(privilege: advanced)# }::
Resumes a failed migrate-root operation if the new_root aggregate is
created and the old root aggregate is in the restricted state.

== Examples

The command in the following example starts the root aggregate migration
on node node1:
```

cluster1::> system node migrate-root -node node1 -disklist 1.11.8,1.11.9,1.11.10,1.11.11,1.11.12 -raid-type raid-dp

```
[[IDef9c0c03cb47bcfac99a3eb5ece16550]]
= system node modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify node attributes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node modify` command sets the attributes of a node.

The owner, location, and asset tag attributes are informational only, and
do not affect the operation of the node or the cluster. The cluster
eligibility attribute marks a node as eligible to participate in a
cluster. The epsilon attribute marks a node as the tie-breaker vote if the
```

```
cluster has an even number of nodes.

Any field of type <text> may be set to any text value. However, if the
value contains spaces or other special characters, you must enter it using
double-quotes as shown in the example below.

Use the xref:{relative_path}system-node-show.html[system node show]
command to display the field values that this command modifies.

== Parameters

``-node {<nodename>|local}`` - Node::
This mandatory parameter specifies which node will have its attributes
modified. The value "local" specifies the current node.

``[-owner <text>]`` - Owner::
This optional text string identifies the node's owner. Fill it in as
needed for your organization.

``[-location <text>]`` - Location::
Use this text string to identify the physical location of the node. This
text string is optional; fill it in as needed for your organization.

``[-assettag <text>]`` - Asset Tag::
If your organization uses asset tags to track equipment, you can use this
text string to store that tag's value.

``[-eligibility {true|false}]`` - Eligibility [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies whether the node is eligible to participate in a
cluster. If you modify another node's eligibility to false, it will no
longer be visible from other nodes in the cluster. If you modify the local
node's eligibility to false, the node will no longer be active in the
cluster and you will not be able to see any cluster nodes from it.

``[-epsilon {true|false}]`` - Epsilon [.advaced-priv]#(privilege:
advanced)#::
If specified as true for a node, this value designates the specified node
as epsilon for this cluster. In a cluster, only one node can be designated
as epsilon at any given time. A node can be designated as Epsilon to add
weight to its voting in a cluster with an even number of nodes.

``[-skip-quorum-check-before-eligible <true>]`` - Skip Quorum Check Before
Setting Node Eligible [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, quorum checks will be skipped prior to
setting a node eligible. When setting a node to eligible, the operation
will continue even if there is a possible data outage due to a quorum
```

```
issue.

``[-skip-quorum-check-before-ineligible <true>]`` - Skip Quorum Check
Before Setting Node Ineligible [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, quorum checks will be skipped prior to
setting a node ineligible. When setting a node to ineligible, the
operation will continue even if there is a possible data outage due to a
quorum issue.

``[-is-diff-svcs {true|false}]`` - Differentiated Services::
If set to ` true ` this means that the specified node and its HA partner
is part of differentiated services storage infrastructure. The default
value for this setting is false.

== Examples

The following example modifies the attributes of a node named node0. The
node's owner is set to "IT" and its location to "Data Center 2."
```

cluster1::> system node modify -node node0 -owner "IT" -location "Data Center 2"

```
== Related Links

* xref:{relative_path}system-node-show.html[system node show]




[[ID8f75452c916b2b7d5fcca495aa81ee27]]
= system node reboot
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reboot a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node reboot` command restarts a node. You can supply a reason
for the reboot, which is stored in the audit log. You can also keep
partner nodes from performing storage takeover during the reboot and
```

```
instruct the rebooted node to create a core dump.

== Parameters

``-node {<nodename>|local}`` - Node::
Specifies the node that is to be restarted. The value "local" specifies
the current node.

``[-f, -inhibit-takeover <true>]`` - Disallow Storage Takeover by
Partner::
If set to `true` , this parameter specifies that the node's failover
partner is not allowed to take over for the node when the node is
rebooted. LIFs will migrate prior to reboot even when inhibit-takeover is
set to `true` . To prevent LIF migration, skip-lif-migration-prior-to-
reboot should be set to `true` . In a two-node MetroCluster configuration,
this parameter prevents automatic unplanned switchover. If you enter this
command without using this parameter, its effective value is false and
storage takeover is allowed. If you enter this parameter without a value,
it is automatically set to true and storage takeover is disabled during
reboot.

``[-reason <text>]`` - Reason for Reboot::
Use this parameter to enter a brief note to indicate the reason for the
restart, which will be stored in the audit log. Providing this information
assists support personnel with troubleshooting efforts.

``[-d, -dump <true>]`` - Create a Core Dump::
If you would like the node to create a core dump before restarting,
specify the `true` value with this parameter. If you enter this command
without using this parameter, its effective value is false and the node
doesn't create a core dump. If you enter this parameter without a value,
it is automatically set to true and the node creates a core dump.

``[-skip-lif-migration-before-reboot <true>]`` - Skip Migrating LIFs Away
from Node Prior to Reboot::
If this parameter is specified, LIF migration prior to the reboot will be
skipped. However if LIFs on this node are configured for failover, those
LIFs may still failover after the reboot has occurred. The default is to
migrate LIFs prior to the reboot. In the default case, the command
attempts to synchronously migrate data and cluster management LIFs away
from the node prior to reboot. If the migration fails or times out, the
reboot will be aborted.

``[-ignore-quorum-warnings <true>]`` - Skip Quorum Check Before Reboot::
If this parameter is specified, quorum checks will be skipped prior to the
reboot. The operation will continue even if there is a possible data
```

```
outage due to a quorum issue.

``[-ignore-strict-sync-warnings <true>]`` - Skip SnapMirror Synchronous
Strict Sync Check Before Reboot::
If this parameter is specified, the check for volumes that are in
SnapMirror Synchronous relationships with policy of type strict-sync-
mirror will be skipped. The operation will continue even if there is a
possible data outage due to not being able to fully sync data.

== Examples

The command in the following example restarts the node named cluster1 for
a software upgrade:
```

cluster1::> system node reboot -node cluster1 -reason "software upgrade"

```
[[ID9b9570cb548577a3041142eb7292bfd2]]
= system node rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node rename` command changes a node's name. Both the node to
be modified and the new name of that node must be specified with the
following parameters. This command is best executed from the node that is
being renamed, using the `-node local` parameter.

Use the xref:{relative_path}system-node-show.html[system node show]
command to display the names of all the nodes in the current cluster.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies which node you are renaming. The value `local`
specifies the current node.

``-newname <text>`` - New Name::
```

```
Use this parameter to specify the new name of the node.
+
  * The name must contain only the following characters: A-Z, a-z, 0-9, "-
" or "_".
  * The first character must be one of the following characters: A-Z or a-
z.
  * The last character must be one of the following characters: A-Z, a-z
or 0-9.
  * The maximum supported length is 47 characters.
  * The system reserves the following names: "all", "cluster", "local" and
"localhost".



== Examples

The following example changes the name of the node named node3 to node4.
```

cluster1::> system node rename -node node3 -newname node4

```
== Related Links

* xref:{relative_path}system-node-show.html[system node show]




[[ID99e559546100d54ade41a9466bbd8c4c]]
= system node restore-backup
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore the original backup configuration to the HA target node

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node restore-backup` command restores the backup configuration
file that is stored on the partner node to the specified target node in an
HA pair. The backup configuration file is restored after ONTAP has been
installed on the target node.
```

```
The backup configuration file is stored on the HA partner node while the
target node is down. After the target node has been installed, the partner
node sends this backup configuration file to the target node through the
management network by using the `system node restore-backup` command to
restore the original configuration. This procedure is commonly used when
replacing the target node's boot device.

The target IP address should be the address of the target node used for
netboot installation.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Specifies the partner node that sends the backup configuration file to the
target node. The value "local" specifies the current node.

``-target-address <Remote InetAddress>`` - HA Partner IP Address
[.advaced-priv]#(privilege: advanced)#::
Specifies the IP address for the target node.




[[ID362eadef1bbcf68add48136238a5bbe4]]
= system node revert-to
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Revert a node to a previous release of Data ONTAP

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node revert-to` command reverts a node's cluster configuration
to the given version. After the `system node revert-to` command has
finished, the `revert_to` command must be run from the nodeshell. The
revert_to command reverts the filesystem on individual nodes to the target
release. Before running `revert-to` in the cluster shell, the target
release must be installed on the node.
```

```
== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
 Specifies the node that is to be reverted. The value `local` specifies
the current node.

``-version <revert version>`` - Data ONTAP Version [.advaced-
priv]#(privilege: advanced)#::
Specifies the version of ONTAP to which the node is to be reverted.

``[-check-only <true>]`` - Capability Check [.advaced-priv]#(privilege:
advanced)#::
If set to `true` , this parameter specifies that the cluster configuration
revert should perform checks to verify all of the preconditions necessary
for revert-to to complete successfully. Setting the parameter to `true`
does not run through the actual revert process. By default this option is
set to `false` .

== Examples

The command in the following example reverts cluster configuration of a
node named node1 to ONTAP version 9.15
```

cluster1::*> system node revert-to -node node1 -version 9.15

```
[[IDbfaffbab187d58911dc61f0eb412893f]]
= system node run-console
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Access the console of a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command allows you to access the console of any remote node on the
same cluster. The remote access is helpful in situations where the node
cannot be booted up or has network issues. This command establishes an SSH
session with the Service Processor of a remote node and accesses that
node's console over the serial channel. This command works even if ONTAP
is not booted up on the remote node. You can get back to the original node
by pressing Ctrl+D. This command works only on SSH sessions and not on
physical console sessions.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node whose physical console you want to
access.

== Examples

The following example accesses the console of node2 in the same cluster.
```

cluster1::> system node run-console -node node2
Pressing Ctrl-D will end this session and any further sessions you might open on top of this session.
Type Ctrl-D.
SP-login: admin
Password:
*
*** This is an SP console session. Output from the ***
*** serial console is also mirrored on this session. ***
*
node2::>
node2::> Connection to 192.168.1.202 closed.

cluster1::>

```
[[IDc34a796b2fd70568abf9236417d3a8a2]]
= system node run
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Run interactive or non-interactive commands in the nodeshell

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Use the `system node run` command to run certain commands from the
nodeshell CLI on a specific node in the cluster. You can run a single
nodeshell command from the clustershell that returns immediately, or you
can start an interactive nodeshell session from which you can run multiple
nodeshell commands.

Nodeshell commands are useful for root volume management and system
troubleshooting. Commands that are available through the nodeshell are
scoped to a single node in the cluster. That is, they affect only the node
specified by the value of the `-node` parameter and do not operate on
other nodes in the cluster. To see a list of available nodeshell commands,
type '?' at the interactive nodeshell prompt. For more information on the
meanings and usage of the available commands, use the `man` command in the
nodeshell.

Only one interactive nodeshell session at a time can be run on a single
node. Up to 24 concurrent, non-interactive sessions can be run at a time
on a node.

When running the nodeshell interactively, exit the nodeshell and return to
the clustershell by using the `exit` command. If the nodeshell does not
respond to commands, terminate the nodeshell process and return to the
clustershell by pressing Ctrl-D.

The `system node run` command is not available from the GUI interface.

NOTE: An alternate way to invoke the `system node run` command is by
typing the `run` as a single word.
```

```
== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the name of the node on which you want to
run the nodeshell command. If you specify only this parameter, the command
starts an interactive nodeshell session that lasts indefinitely. You can
exit the nodeshell to the clustershell by pressing Ctrl-D or by typing the
`exit` command.

{ ``[-command <text>,...]`` - Command to Run::
This optionally specifies the name of a single nodeshell command to run on
the specified node. To see a list of available nodeshell commands, type
'?' at an interactive nodeshell prompt.

| ``[-reset <true>]`` - Reset Existing Connection }::
If this parameter is specified with the `true` value, it terminates any
existing interactive nodeshell session on the specified node. The default
value is `false` .

== Examples

The following example runs the nodeshell command `sysconfig -V` on a node
named node1:
```

cluster1::> system node run -node node1 -command sysconfig -V
volume node1_aggr0 (1 RAID group):
group 0: 3 disks

```
The following example starts a nodeshell session on a node named node2 and
then runs the nodeshell `sysconfig -V` command. The system remains in the
nodeshell after running the `sysconfig -V` command.
```

cluster1::> run -node node2
Type 'exit' or 'Ctrl-D' to return to the CLI
node2> sysconfig -V
volume node2_aggr0 (1 RAID group):
group 0: 3 disks
node2>

```
The following example starts a nodeshell session on a node named node1 and
then runs two nodeshell commands, `aggr status` first and `vol status`
second. Use quotation marks and semicolons when executing multiple
nodeshell commands with a single `run` command.
```

cluster1::> run -node node1 -command "aggr status; vol status"

Aggr State Status Options
aggr0 online raid_dp, aggr root
parity uninit'd!
32-bit
aggr1 online raid_dp, aggr
parity uninit'd!
32-bit
Volume State Status Options
vol0 online raid_dp, flex root, nvfail=on
parity uninit'd!
root_vs0 online raid_dp, flex create_ucode=on,
cluster convert_ucode=on,
parity uninit'd! maxdirsize=102400

```
[[IDda9aa65cde36aed2a08b19a3cda85849]]
= system node show-discovered
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display all nodes discovered on the local network

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node show-discovered` command displays information about all
the detectable nodes on the local cluster network. This includes both
nodes in a cluster and nodes that do not belong to a cluster. You can
filter the output to show only nodes that do not belong to a cluster or
nodes that are in a cluster.

To see a list of values that are in use for a particular field, use the `-
fields` parameter of this command with the list of field names you want to
view.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If the `-instance` parameter is specified, the command displays detailed
information about each node.
```

```
``[-node <text>]`` - Node Name::
This parameter specifies the name of a node for which information is to be
displayed. If this parameter is not specified, the command displays
information about all discovered nodes.

``[-is-in-cluster {true|false}]`` - Is in a Cluster::
If this parameter is set to `false` , the command lists only nodes that do
not belong to a cluster.

``[-cluster-uuid <UUID>]`` - Cluster UUID::
Displays information about nodes belonging to the cluster that has the
UUID you specify.

``[-cluster-name <text>]`` - Cluster Name::
Displays information about nodes belonging to the cluster that has the
name you specify.

``[-serial-number <text>]`` - Node Serial Number::
Displays information about the node that has the serial number you
specify.

``[-addresses <IP Address>,...]`` - Cluster IP Addresses::
Displays information about the node that has the cluster IP addresses you
specify.

``[-netmask <IP Address>]`` - Cluster Address Mask::
Displays information about the nodes that have the netmask address you
specify.

``[-nvramid <nvramid>]`` - Node NVRAM ID::
Displays information about the node that has the NVRAM ID you specify.

``[-partner-nvramid <nvramid>]`` - Partner NVRAM ID::
Displays information about the node that has an HA partner with the NVRAM
ID you specify.

``[-model <text>]`` - Model::
Displays the nodes that have the specified model number.

``[-version <text>]`` - Software Version::
Displays the nodes that have the specified version of ONTAP.

== Examples

The following example displays information about all discovered nodes in
the cluster network:
```

```
cluster1::*> system node show-discovered
Node Cluster Addresses NVRAM ID Partner NVRAM
------------- ------------- --------------- ----------- --------------
4069114-60-0 - 169.254.232.178 4069114600 -
4069114-60-2 - 169.254.79.38 4069114602 -
4069114-60-3 - 169.254.195.76 4069114603 -
cluster1-01 cluster1 169.254.140.39 4069114628 4069114629
cluster1-02 cluster1 169.254.138.137 4069114629 4069114628
```

```
[[IDe401c080c5fdd847d4c15819e3185a4f]]
= system node show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of nodes in the cluster

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node show` command displays information about the nodes in a
cluster. You can limit output to specific types of information and
specific nodes in the cluster, or filter output by specific field values.

To see a list of values that are in use for a particular field, use the `-
fields` parameter of this command with the list of field names you want to
view. Use the xref:{relative_path}system-node-modify.html[system node
modify] command to change some of the field values that this command
displays.

== 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.

| ``[-inventory  ]``::
Use this parameter to display inventory information such as serial
numbers, asset tags, system identifiers, and model numbers.

| ``[-messages  ]``::
Use this parameter to display system messages for each node.
```

```
| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Selects information for node names that match this parameter value.

``[-owner <text>]`` - Owner::
Selects nodes that have owner values that match this parameter value.

``[-location <text>]`` - Location::
Selects nodes at specific physical locations that match this parameter
value.

``[-model <text>]`` - Model::
Selects nodes that have model numbers that match this parameter value.

``[-serialnumber <text>]`` - Serial Number::
Selects nodes that have serial numbers that match this parameter value.

``[-assettag <text>]`` - Asset Tag::
Selects nodes that have asset tags that match this parameter value.

``[-uptime {<seconds>|[<d> days] <hh>:<mm>[:<ss>]}]`` - Uptime::
Selects nodes that have uptimes that match this parameter value. This
parameter is most useful when used with a range indicator such as less
than or greater than, as in:
```

```
show -uptime >"30 days 00:00"
```

```
``[-nvramid <nvramid>]`` - NVRAM System ID::
Selects nodes that have NVRAM system IDs that match this parameter value.

``[-systemid <text>]`` - System ID::
Selects nodes that have system IDs that match this parameter value.

``[-vendor <text>]`` - Vendor::
Selects nodes that have vendor names that match this parameter value.

``[-health {true|false}]`` - Health::
Selects nodes that have health values that match this parameter value.
Specify `true` to display healthy nodes, and `false` to display unhealthy
nodes.
```

```
``[-eligibility {true|false}]`` - Eligibility::
Selects nodes that have voting eligibility values that match this
parameter value.

``[-epsilon {true|false}]`` - Epsilon [.advaced-priv]#(privilege:
advanced)#::
Selects nodes that have epsilon holding designations that match this
parameter value. This is useful to find out which node, if any, in the
current cluster has been designated as epsilon. Specify `true` to display
the node holding epsilon, and `false` to display nodes not holding
epsilon.

``[-uuid <UUID>]`` - UUID [.advaced-priv]#(privilege: advanced)#::
Selects nodes that have the specified universal unique identifiers that
match this parameter value.

``[-is-diff-svcs {true|false}]`` - Differentiated Services::
If true, the corresponding node is considered to be part of differentiated
services storage infrastructure.

``[-is-all-flash-optimized {true|false}]`` - All-Flash Optimized::
Selects nodes that have "All-Flash Optimized" personality values that
match this parameter value. Specify ``_true_`` to display nodes which
support only SSD drives, and ``_false_`` to display nodes which support
all kinds of drives.

``[-is-capacity-optimized {true|false}]`` - Capacity Optimized::
Selects nodes that have "Capacity Optimized" personality values that match
this parameter value. Specify ``_true_`` to display nodes which support
only SSD drives with Capacity Optimized personality enabled and set
``_false_`` otherwise.

``[-is-qlc-optimized {true|false}]`` - FAS QLC Optimized::
Selects nodes that have "QLC Optimized" personality values that match this
parameter value. Specify ``_true_`` to display nodes which support only
SSD drives with QLC Optimized personality enabled and set ``_false_``
otherwise.

``[-is-all-flash-select-optimized {true|false}]`` - All-Flash Select
Optimized::
Selects nodes that have "All-Flash Select Optimized" personality values
that match this parameter value. Specify ``_true_`` to display nodes which
support only SSD drives, and ``_false_`` to display nodes which support
all kinds of drives.
```

```
``[-sas2-sas3-mixed-stack-support {all|direct-attached|bridge-
attached|none}]`` - SAS2/SAS3 Mixed Stack Support::
Selects nodes that have "SAS2/SAS3 Mixed Stack Support" values that match
this parameter value. The possible values are:
+
  * ``_all_`` : SAS2/SAS3 mixing supported on all stacks
  * ``_direct-attached_`` : SAS2/SAS3 mixing supported on direct-attached
stacks
  * ``_bridge-attached_`` : SAS2/SAS3 mixing supported on bridge-attached
stacks
  * ``_none_`` : SAS2/SAS3 mixing not supported



== Examples

The following example displays information about all nodes in the cluster:
```

cluster1::> system node show
Node Health Eligibility Uptime Model Owner Location
------ ------ ----------- ------------- -------- ------ -------------
node0 true true 89 days 23:47 MODELXX IT Data Center 2
node1 true true 15 days 22:37 MODELXX Data Center 2
node2 true true 15 days 23:00 MODELXX Data Center 2
node3 true true 15 days 22:37 MODELXX Data Center 2
4 entries were displayed.

```
This example displays the locations and model numbers of all nodes that
are in physical locations that have names beginning with "Lab":
```

cluster1::> system node show -location lab* -fields location, model
node location model
------------- -------- ------
node5 Lab 1 MODELXX
node7 Lab 3 MODELXX
node9 Lab 5 MODELXX

```
== Related Links

* xref:{relative_path}system-node-modify.html[system node modify]




[[ID5be136fbddeb874aa398ffce8c4f25da]]
= system node autosupport invoke-core-upload
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Generate and send an AutoSupport message with an existing core file.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport invoke-core-upload` command sends two
AutoSupport messages. The first AutoSupport message contains the content
relevant to this core upload. This AutoSupport message has subject line
with prefix "CORE INFO:". The second Autosupport message contains the core
file specified by the "-core-filename" option. This AutoSupport message
has subject line with prefix "CORE UPLOAD:". The command requires that the
specified file be present while the AutoSupport message is being
transmitted.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node from which the AutoSupport message
is sent. Defaults to localhost.

``[-message <text>]`` - Message Included in the AutoSupport Subject::
Use this parameter to specify the text in the subject line of the
AutoSupport message.

``[-uri <text>]`` - Alternate Destination for This AutoSupport::
Use this parameter to send the AutoSupport message to an alternate
destination. Only "https" protocol is supported. If this parameter is
omitted, the message is sent to all the recipients defined by the
xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify] command.

``[-force <true>]`` - Generate and Send Even if Disabled::
Use this parameter to generate and send the AutoSupport message even if
AutoSupport is disabled on the node.

``[-case-number <text>]`` - Case Number for This Core Upload::
Use this parameter to specify the optional case number to be associated
with this AutoSupport message.
```

```
``-core-filename <text>`` - The Existing Core Filename to Upload::
Use this parameter to specify the core file to be included in the
AutoSupport message. Use the xref:{relative_path}system-node-coredump-
show.html[system node coredump show] command to list the core files by
name.

== Examples

Use this command to list the core files from a node:
```

cluster1::> system node coredump show -node local
Node:Type Core Name Saved Panic Time
--------- ---------- ----- ----------
node:kernel
core.4073000068.2013-09-11.15_05_01.nz true 9/11/2013 15:05:01

```
Use this command to invoke an AutoSupport message with the corefile
core.4073000068.2013-09-11.15_05_01.nz:
```

cluster1::> system node autosupport invoke-core-upload -core-filename core.4073000068.2013-09-11.15_05_01.nz -node local

```
== Related Links

* xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify]
* xref:{relative_path}system-node-coredump-show.html[system node coredump
show]




[[IDba4377b021e6c083ea56cbc1d998ea68]]
= system node autosupport invoke-diagnostic
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Generate and send an AutoSupport message with diagnostic content from
specified subsystems.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.
```

```
== Description

The `system node autosupport invoke-diagnostic` command sends an
AutoSupport message from a node containing basic content from all
subsystems along with troubleshooting and diagnostic content from
specified subsystems.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node from which the AutoSupport message
is sent.

``[-message-subject <text>]`` - Message Included in the AutoSupport
Subject [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify text sent in the subject line of the
AutoSupport message.

``[-uri <text>]`` - Alternate Destination for This AutoSupport [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to send the AutoSupport message to the destination you
specify instead of the configured destination. Only "file", "mailto" and
"https" protocols are supported. If this parameter is omitted, the message
is sent to the all of the recipients defined by the
xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify] command.

``[-force <true>]`` - Generate and Send Even if Disabled [.advaced-
priv]#(privilege: advanced)#::
This flag indicates this AutoSupport should be generated and delivered
even if the AutoSupport configuration is disabled. If "0", then obey the
normal "state", "support" and "diagnostic-content-state" flags. If this is
"1", then ignore the "state", "support" and "diagnostic-content-state".

``-subsystems <subsys1,subsys2,...>,...`` - Subsystems to Collect Basic,
Troubleshooting & Diagnostic Content [.advaced-priv]#(privilege:
advanced)#::
Use this mandatory parameter to include basic, troubleshooting, and
diagnostic configuration content from the subsystems you specify. Content
is collected from the "mandatory" subsystem and the subsystems you
specify. You can specify up to four subsystems.

== Examples

The following example sends a diagnostic AutoSupport message from a node
```

```
named myNode containing information from the "mandatory" subsystem and
basic, troubleshooting and diagnostic contents from "raid", "wafl" and
"ha" subsystems:
```

cluster1::> system node autosupport invoke-diagnostic -node myNode -subsystems raid,wafl,ha

```
== Related Links

* xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify]




[[ID6086a1ed855fe4ae2eb363627ffe75a9]]
= system node autosupport invoke-performance-archive
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Generate and send an AutoSupport message with performance archives.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport invoke-performance-archive` command sends an
AutoSupport message with the performance archives from a node. The command
requires that the performance archives in the specified date range be
present while the AutoSupport message is being transmitted.

NOTE: Performance archive upload is not supported when the `-remove-
private-data` parameter is enabled. To disable this parameter, use the
xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify -remove-private-data false] command.



== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node from which the AutoSupport message
is sent. The default setting is localhost.
```

```
``[-message <text>]`` - Message Included in the AutoSupport Subject::
Use this parameter to specify the text in the subject line of the
AutoSupport message.

``[-uri <text>]`` - Alternate Destination for This AutoSupport::
Use this parameter to send the AutoSupport message to an alternate
destination. Only "file" and "https" protocols are supported. If this
parameter is omitted, the message is sent to the all of the recipients
defined by the system node autosupport modify command.

``[-force <true>]`` - Generate and Send Even if Disabled::
Use this parameter to generate and send the AutoSupport message even if
AutoSupport is disabled on the node.

``[-case-number <text>]`` - Case Number for This Performance Archive
Upload::
Use this parameter to specify the optional case number to be associated
with this AutoSupport message.

``-start-date <MM/DD/YYYY HH:MM:SS>`` - Start Date for Performance Archive
Dataset::
Use this parameter to specify the start date for the files in the
performance archive dataset to be included in the AutoSupport message.

{ ``-end-date <MM/DD/YYYY HH:MM:SS>`` - End Date for Performance Archive
Dataset::
Use this parameter to specify the end date for the files in the
performance archive dataset to be included in the AutoSupport message. The
end date should be within six hours of the start date.

| ``-duration <[<integer>h][<integer>m][<integer>s]>`` - Duration of
Performance Archive Dataset }::
Use this parameter with start-date to specify the duration of the
performance archive dataset to be included in the AutoSupport message. The
maximum duration limit is six hours from the start date.

== Examples

Use this command to invoke an AutoSupport message to include the
performance archives in the given date range:
```

cluster1::> system node autosupport invoke-performance-archive -node local -start-date 9/12/2013 19:37:24 -end-date 9/12/2013 22:37:24
cluster1::> system node autosupport invoke-performance-archive -node local -start-date 11/21/2013 13:42:09 -duration 6h

```
== Related Links

* xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify]




[[ID436486fe6c627e135fe2282e0bf9d7f5]]
= system node autosupport invoke-splog
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Generate and send an AutoSupport message with collected service-processor
log files

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport invoke-splog` command sends an AutoSupport
message with the Service Processor (SP) or baseboard management controller
(BMC) log files from a specified node in the cluster.

== Parameters

``-remote-node {<nodename>|local}`` - Node::
Use this parameter to specify the node from which Service Processor log
files are to be collected.

``[-log-sequence <integer>]`` - Log File Sequence Number::
Use this parameter to specify the sequence number of the Service Processor
log files to be collected. If this parameter is omitted, the latest
Service Processor log files are collected.

``[-uri <text>]`` - Alternate Destination for This AutoSupport::
Use this parameter to send the AutoSupport message to an alternate
destination. Only "file" and "https" protocols are supported. If this
parameter is omitted, the message is sent to the all of the recipients
defined by the xref:{relative_path}system-node-autosupport-
modify.html[system node autosupport modify] command.

``[-force <true>]`` - Generate and Send Even if Disabled::
```

```
Use this parameter to generate and send the AutoSupport message even if
AutoSupport is disabled on the node.

``[-full <true>]`` - Full Log Collection::
Use this parameter to collect full logs from the SP or BMC.

``[-show-full-progress <true>]`` - Full Log Collection Progress Status::
Use this parameter to find the status of a full log collection from a SP
or BMC.

== Examples

Use this command to invoke an AutoSupport message to include the Service
Processor log files collected from node cluster1-02.
```

cluster1::> system node autosupport invoke-splog -remote-node cluster1-02
[Job 777] Job succeeded: Log files from the service processor have been transferred to "/mroot/etc/log/sp/ondemand" on node cluster1-01, and AutoSupport message has been triggered.

cluster1::>

```
== Related Links

* xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify]




[[ID9abca0a8841cf72f44a2dc2d00bdd908]]
= system node autosupport invoke
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Generate and send an AutoSupport message

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport invoke` command sends an AutoSupport message
```

```
from a node.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node from which the AutoSupport message
is sent.

``[-message <text>]`` - Message Included in the AutoSupport Subject::
Use this parameter to specify text sent in the subject line of the
AutoSupport message. This parameter is not available when the `-type`
parameter is set to performance.

``-type {test|performance|all}`` - Type of AutoSupport Collection to
Issue::
Use this parameter to specify the type of AutoSupport collection to issue.
There is no default; you must specify a `-type` .
+
  * test - The message contains basic information about the node. When the
AutoSupport message is received by technical support, an e-mail
confirmation is sent to the system owner of record. This enables you to
confirm that the message is being received by technical support.
  * all - The message contains all collected information about the node.
  * performance - The message contains only performance information about
the node. This parameter has effect only if performance AutoSupport
messages are enabled, which is controlled by the `-perf` parameter of the
xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify] command.



``[-uri <text>]`` - Alternate Destination for This AutoSupport::
Use this parameter to send the AutoSupport message to the destination you
specify instead of the configured destination. Only "file", "mailto" and
"https" protocols are supported. If this parameter is omitted, the message
is sent to the all of the recipients defined by the
xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify] command.

``[-force <true>]`` - Generate and Send Even if Disabled::
Use this parameter to generate and send the message even if AutoSupport is
disabled on the node.

== Examples

The following example sends a test AutoSupport message from a node named
node0 with the text "Testing ASUP":
```

```
cluster1::> system node autosupport invoke -node node0 -type test -message "Testing ASUP"
```

```
== Related Links

* xref:{relative_path}system-node-autosupport-modify.html[system node
autosupport modify]




[[ID3b7fc1ea4ba0bf3df7d1671aeaad07b4]]
= system node autosupport modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify AutoSupport configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport modify` command modifies the AutoSupport
configuration of all the nodes in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node::
+
NOTE: The AutoSupport configuration will be modified on all nodes in the
cluster, even if this parameter is specified.
+
This parameter is ignored, but retained for CLI backward compatibility.

``[-state {enable|disable}]`` - State::
Use this parameter to specify whether AutoSupport is enabled or disabled.
The default setting is `enable` . When AutoSupport is disabled, messages
are not sent to anyone, including the vendor's technical support, your
internal support organization, or partners.

``[-mail-hosts <text>,...]`` - SMTP Mail Hosts::
Use this parameter to specify up to five SMTP mail hosts through which
AutoSupport messages are sent out. This parameter is required if you
```

```
specify e-mail addresses in the -to, -noteto, or -partner-address
parameters or if you specify smtp in the -transport parameter. Separate
multiple mail hosts with commas and do not use spaces in between. The
AutoSupport delivery engine attempts to use these hosts for delivery in
the order that you specify.
+
You can optionally specify a port value for each mail server. A port value
can be specified on none, all, or some of the mail hosts. The port
specification for a mail host consists of a colon (":") and a decimal
value between 1 and 65335, and follows the mailhost name (for example,
``_mymailhost.example.com:5678_`` ). The default port value is ``_587_``
if `-smtp-encryption` is set to ``_start_tls_`` . Otherwise, the default
is ``_25_`` .
+
Also, you can optionally specify a user name and password combination for
authentication to each mail server. The format when specifying a username
is ``_user1@mymailhost.example.com_`` . User will be prompted for the
password. The username and password can be specified on none, all, or some
of the mail hosts.
+
If the user name itself includes the "@" symbol and domain information,
the format to be used is ``_user%40domain.com@mymailhost.example.com_`` .
User will be prompted for the password.
+
The default value for this parameter is ``_mailhost_`` .

``[-from <mail address>]`` - From Address::
Use this parameter to specify the e-mail address from which all the nodes
in the cluster send AutoSupport messages. The default is Postmaster. If
you require node-specific "from" addresses, enable the node-specific-from
(privilege: advanced) parameter.

``[-to <mail address>,...]`` - List of To Addresses::
Use this parameter to specify up to five e-mail addresses to receive
AutoSupport messages that are most relevant for your internal
organization. Separate multiple addresses with commas with no spaces in
between. For this parameter to have effect, the -mail-hosts parameter must
be configured correctly. Individual trigger events can change the default
usage of this parameter using the `-to` parameter of the
xref:{relative_path}system-node-autosupport-trigger-modify.html[system
node autosupport trigger modify] command. By default, no list is defined.

``[-noteto <mail address>,...]`` - (DEPRECATED) List of Noteto Addresses::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP. Instead, use the xref:{relative_path}event-notification-
```

```
destination-create.html[event notification destination create] command to
create an email destination and the xref:{relative_path}event-
notification-create.html[event notification create] command to map the
`important-events` system filter to the newly created event destination.
This will send an email notification for all events with a severity of
ERROR, ALERT, or EMERGENCY. For events with other severity values, a new
event filter needs to be created using the xref:{relative_path}event-
filter-create.html[event filter create] command. Filter rules can be added
to the filter using the xref:{relative_path}event-filter-rule-
add.html[event filter rule add] command. This newly created filter has to
be mapped to the event destination.
+
Use this parameter to specify up to five addresses to receive a short-note
version of AutoSupport messages that are most relevant for your internal
organization. Short-note e-mails contain only the subject line of the
AutoSupport message, which is easier to view on a mobile device. For this
parameter to have effect, the `-mail-hosts` parameter must be configured
correctly. Individual trigger events can change the default usage of this
parameter using the `-noteto` parameter of the xref:{relative_path}system-
node-autosupport-trigger-modify.html[system node autosupport trigger
modify] command. By default, no list is defined.

``[-partner-address <mail address>,...]`` - List of Partner Addresses::
Use this parameter to specify up to five e-mail addresses to receive all
AutoSupport messages including periodic messages. This parameter is
typically used for support partners. For this parameter to have effect,
the -mail-hosts parameter must be configured correctly. By default, no
list is defined.

``[-support {enable|disable}]`` - Send AutoSupport Messages to Vendor
Support::
Use this parameter to specify whether to send all AutoSupport messages to
your vendor's technical support. (Destination information is pre-defined
and does not require configuration.) When -state is enabled and -support
is disabled, messages are sent to the addresses specified in the -to,
-noteto, or -partner-address parameters but are not sent to your vendor's
technical support. The default is `enable` .

``[-transport {smtp|https}]`` - Protocol to Contact Support::
+
NOTE: ``_http_`` transport is no longer supported by AutoSupport servers.
+
Use this parameter to specify the protocol used to deliver AutoSupport
messages to your vendor's technical support. This parameter applies only
when the -support parameter is set to enable. If you specify https and
your network uses a proxy, you must also set the -proxy-url parameter. If
```

```
you specify smtp, you must also configure the -mail-hosts parameter.

``[-proxy-url <text>]`` - Support Proxy URL::
Use this parameter to specify an HTTP/S proxy if the `-transport`
parameter is set to HTTPS and your organization uses a proxy. Enter the
URL with an http:// or https:// prefix. If a scheme is not specified with
the URL, it is assumed to be http://. If authentication is required, use
the format "[username]@[host][:[port]]". You will be prompted for the
password. If using an HTTPS proxy, you also need to install the correct
Root CA certificates in ONTAP. The default is an empty string. To specify
a proxy that contains a question mark, press ESC followed by the "?". This
field can be cleared by setting the value to an empty string using two
double quotes ("").

``[-hostname-subj {true|false}]`` - Hostname Subject::
Use this parameter to specify whether the hostname of the node is included
in the subject line of the AutoSupport message. The default is `false` .
This parameter applies only if the `-remove-private-data` parameter is
`true` .

``[-perf {true|false}]`` - Performance Data Enable::
Use this parameter to specify whether performance data is sent to
technical support and addresses specified in the -partner-address
parameter. The default is `true` .

``[-retry-interval <[<integer>h][<integer>m][<integer>s]>]`` - Retry
Interval::
Use this parameter to specify the amount of time to delay before trying to
send an AutoSupport message again after a sending failure. Values may end
with "s", "m", or "h" to indicate seconds, minutes, or hours,
respectively. The minimum interval is 30 seconds and the maximum is 1 day.
The default is 4 minutes.

``[-retry-count <integer>]`` - Retry Count::
Use this parameter to specify the number of times to try resending mail
before dropping it. The minimum number is 5 and the maximum is 30. The
default is 15 times.

``[-reminder {true|false}]`` - Reminder Enable::
Use this parameter to enable or disable a reminder message that is sent
when AutoSupport is not configured to send messages to technical support.
This reminder is logged as an EMS event called
"autosupport.general.reminder" every 24 hours. The default is `true` .

``[-max-http-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum HTTPS Size::
Use this parameter to specify the maximum file size (in bytes by default,
```

```
but can also be specified in KB, MB, TB or PB) for HTTPS transfers. This
parameter applies only to messages sent to technical support and only if
the -transport parameter is set to HTTPS. Setting the value to 0 disables
the delivery size budget. The default is 50 MB and the minimum supported
size is 2 MB.
+
If the size of the AutoSupport message exceeds this value, AutoSupport
will deliver as much of the message as possible. You can use the "system
node autosupport manifest show" command to identify the sections of the
message that AutoSupport sent. AutoSupport collects and sends the content
in order of priority. The priority is predefined for each AutoSupport
message. To identify the collection order for an AutoSupport trigger, use
the "system node autosupport trigger show" command with the -instance
parameter.

``[-max-smtp-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum SMTP Size::
Use this parameter to specify the maximum file size (in bytes by default,
but can also be specified in KB, MB, TB or PB) for SMTP (e-mail)
transfers. This parameter applies to messages sent to the addresses
specified in the -to, -noteto, and -partner-address parameters. If the
-transport parameter is set to smtp, this parameter also applies to
messages sent to the vendor's technical support. Setting the value to 0
disables the delivery size budget. The default is 5 MB and the minimum
supported size is 2 MB.
+
If the size of the AutoSupport message exceeds this value, AutoSupport
will deliver as much of the message as possible. You can use the "system
node autosupport manifest show" command to identify the sections of the
message that AutoSupport sent. AutoSupport collects and sends the content
in order of priority. The priority is predefined for each AutoSupport
message. To identify the collection order for an AutoSupport trigger, use
the "system node autosupport trigger show" command with the -instance
parameter.

``[-remove-private-data {true|false}]`` - Remove Sensitive Data::
Use this parameter with the value `true` to remove, encode, or mask
sensitive data from AutoSupport attachments and headers. Use this feature
to eliminate private data from all AutoSupport messages. The default is
`false` .
+
Eliminated data can include: IP addresses, MAC addresses, URIs, DNS names,
e-mail addresses, port numbers, node names, Vserver names, cluster names,
aggregate names, volume names, junction paths, policy names, user IDs,
group IDs, LUNs, NVMe namespaces and qtree names.
+
Unstructured log files (such as the EMS event log) are also no longer
```

```
collected by AutoSupport to ensure sensitive data is not included. As a
result, log files will need to be manually collected and provided to
technical support for troubleshooting purposes.
+
NOTE: Changing this value from false to true deletes the AutoSupport
history and all files associated with it.
+


``[-validate-digital-certificate {true|false}]`` - Validate Digital
Certificate Received [.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value `true` to force each node to validate
digital certificates that it receives. The default is `true` . When this
value is true, the certificate might be validated by OCSP. The OCSP
validation for these certificates is controlled by
xref:{relative_path}security-config-ocsp-enable.html[security config ocsp
enable -app autosupport] and xref:{relative_path}security-config-ocsp-
disable.html[security config ocsp disable -app autosupport] .

``[-ondemand-state {enable|disable}]`` - AutoSupport OnDemand State
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify whether the AutoSupport OnDemand feature is
enabled or disabled. The default is `enable` . When AutoSupport OnDemand
is enabled, support personnel can remotely trigger new AutoSupport
messages, re-send existing AutoSupport messages and decline the delivery
of unwanted AutoSupport messages. When this option is disabled, this node
will not respond to any AutoSupport OnDemand requests from support
personnel.

``[-ondemand-remote-diagnostics-state {enable|disable}]`` - AutoSupport
OnDemand Remote Diagnostics State [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify whether the AutoSupport OnDemand Remote
Diagnostics feature is enabled or disabled. The default is `enable` . When
AutoSupport OnDemand Remote Diagnostics is enabled, support personnel can
remotely trigger new AutoSupport messages on this node to gather
information during troubleshooting scenarios. When this option is
disabled, support personnel will still be able to re-send existing
AutoSupport messages that may not have been transmitted correctly.

``[-compress-spool-dir {true|false}]`` - Compress AutoSupport Spool Dir
[.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value `true` to compress the AutoSupport spool
directory once the AutoSupport is in a terminal state. The default is
`true` .

``[-node-specific-from {enable|disable}]`` - Node-Specific From Address
```

```
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify whether the "from" e-mail address needs to
be node-specific. If enabled, the node-specific "from" e-mail address is
generated by prepending the node name and "-" to the configured "from"
address. For example, if the node name is "cluster-01" and the "from"
parameter is "prod@company.com", the generated "from" e-mail address is
"cluster-01-prod@company.com".

``[-validate-smtp-certificate {true|false}]`` - Validate Certificate of
SMTP Server [.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value `true` to force each node to validate
digital certificates that it receives from the SMTP server. The default is
`true` . When this value is true, the certificate might be validated by
OCSP. The OCSP validation for these certificates is controlled by
xref:{relative_path}security-config-ocsp-enable.html[security config ocsp
enable -app autosupport] and xref:{relative_path}security-config-ocsp-
disable.html[security config ocsp disable -app autosupport] .

``[-smtp-encryption {none|start_tls}]`` - SMTP Encryption::
Use this parameter to specify the encryption protocol used to deliver
AutoSupport messages to the configured mail servers. When this parameter
is set to `none` , the SMTP AutoSupport message is not encrypted. When
this parameter is set to `start_tls` , the existing plaintext channel is
upgraded to a TLS encrypted channel using the STARTTLS protocol. If the
TLS connection fails, there is no fallback to plaintext transfer. The
default is `none` .

``[-force <true>]`` - Bypass Interdependency Checks [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to modify some of the AutoSupport configurations that
are otherwise blocked when the automatic update feature is enabled.
Without this parameter, an attempt to disable AutoSupport, modify the
transport to SMTP, or disable the AutoSupport OnDemand feature fails if
the automatic update feature is enabled.

== Examples

The following example enables AutoSupport on all nodes in the cluster with
the following settings:



* SMTP mail host named smtp.example.com.
* E-mail "from" address of alerts@example.com
* E-mail "to" address of support@example.com
* AutoSupport messages sent to support personnel
```

```
* HTTPS set as transport protocol
* If sending fails, the system will wait 23 minutes before retrying.
```

cluster1::> system node autosupport modify -state enable -mail-hosts smtp.example.com -from alerts@example.com -to support@example.com -support enable -transport https -retry-interval 23m

```
The following examples show how to modify AutoSupport URLs when using IPv6
address literals. You will be prompted for the password in the second
example:
```

cluster1::> system node autosupport modify -mail-hosts [2620:10a:4002:6004::bbbb]:25

```

```

cluster1::> system node autosupport modify -proxy-url username@[2620:10a:4002:6004::bbbb]:8080

```
== Related Links

* xref:{relative_path}system-node-autosupport-trigger-modify.html[system
node autosupport trigger modify]
* xref:{relative_path}event-notification-destination-create.html[event
notification destination create]
* xref:{relative_path}event-notification-create.html[event notification
create]
* xref:{relative_path}event-filter-create.html[event filter create]
* xref:{relative_path}event-filter-rule-add.html[event filter rule add]
* xref:{relative_path}security-config-ocsp-enable.html[security config
ocsp enable]
* xref:{relative_path}security-config-ocsp-disable.html[security config
ocsp disable]




[[ID9237afe43fcdc6119a9543ad9c23f0ae]]
= system node autosupport show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display AutoSupport configuration

*Availability:* This command is available to _cluster_ administrators at
```

```
the _admin_ privilege level.

== Description

The `system node autosupport show` command displays the AutoSupport
configuration of one or more nodes.

== 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.

| ``[-config  ]``::
Use this parameter to display the retry interval, retry count, throttle,
and reminder settings of all nodes in the cluster.

| ``[-recent  ]``::
Use this parameter to display the subject and time of the last AutoSupport
message generated by each node in the cluster.

| ``[-support-http  ]``::
Use this parameter to display whether HTTP support is enabled in the
cluster, and identify the transport protocol and the support proxy URL
used.

| ``[-support-smtp  ]``::
Use this parameter to display whether SMTP (e-mail) support is enabled in
the cluster, and identify the transport protocol and the "to" e-mail
address used.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
+
NOTE: The AutoSupport configuration is identical for all nodes in the
cluster.
+
Use this parameter to display detailed information about the node you
specify.

``[-state {enable|disable}]`` - State::
Use this parameter to display information only about nodes that have the
AutoSupport state you specify.
```

```
``[-mail-hosts <text>,...]`` - SMTP Mail Hosts::
Use this parameter to display information only about nodes that use the
mail hosts you specify.
+
The port number used is displayed with the mailhost after a ":". The
default port value is ``_587_`` if `-smtp-encryption` is set to
``_start_tls_`` . Otherwise, the default is ``_25_`` .

``[-from <mail address>]`` - From Address::
Use this parameter to display information only about nodes that have the
"from" e-mail address you specify.

``[-to <mail address>,...]`` - List of To Addresses::
Use this parameter to display information only about nodes that have the
"to" e-mail addresses you specify.

``[-noteto <mail address>,...]`` - (DEPRECATED) List of Noteto Addresses::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP.
+
Use this parameter to display information only about nodes that send
short-note e-mail messages to the e-mail addresses you specify. Short-note
e-mails contain only the subject line of the AutoSupport message, which is
easier to view on a mobile device.

``[-partner-address <mail address>,...]`` - List of Partner Addresses::
Use this parameter to display information only about nodes that have the
"partner-address" e-mail addresses you specify. These addresses are not
subject to the delivery limitations configured for the "-to" addresses of
AutoSupport triggers.

``[-support {enable|disable}]`` - Send AutoSupport Messages to Vendor
Support::
Use this parameter with the value "enable" to display information only
about nodes that send AutoSupport messages to your vendor's technical
support. Use this parameter with the value "disable" to display
information only about nodes that do not send AutoSupport messages to your
vendor's technical support.

``[-transport {smtp|https}]`` - Protocol to Contact Support::
Use this parameter to display information only about nodes that use the
protocol you specify to send AutoSupport messages.

``[-url <text>]`` - Support URL for HTTPS::
```

```
Use this parameter to display information only about nodes that use the
URL you specify to send messages through the HTTPS POST operations.

``[-put-url <text>]`` - Support URL for HTTPS PUT::
Use this parameter to display information only about nodes that use the
URL you specify to send messages through the HTTPS PUT operations.

``[-proxy-url <text>]`` - Support Proxy URL::
Use this parameter to display information only about nodes that use the
proxy URL you specify.

``[-support-address <mail address>,...]`` - Support Address::
Use this parameter to display information only about nodes that use the
external support address you specify.

``[-hostname-subj {true|false}]`` - Hostname Subject::
Use this parameter to display information only about nodes that include
their hostname in the "Subject:" line of AutoSupport messages. If the
parameter "remove-private-data" is `false` , this parameter has no effect.

``[-perf {true|false}]`` - Performance Data Enable::
Use this parameter with the value "true" to display information only about
nodes that send periodic performance AutoSupport messages. Use this
parameter with the value "false" to display information only about nodes
that do not send periodic performance messages.

``[-retry-interval <[<integer>h][<integer>m][<integer>s]>]`` - Retry
Interval::
Use this parameter to display information only about nodes that use the
retry interval you specify.

``[-retry-count <integer>]`` - Retry Count::
Use this parameter to display information only about nodes that use the
retry count you specify.

``[-reminder {true|false}]`` - Reminder Enable::
Use this parameter with the value "true" to display information only about
nodes that send messages reminding administrators to enable AutoSupport if
AutoSupport is not enabled. Use this parameter with the value "false" to
display information only about nodes that do not send reminder messages.

``[-last-subject <text>]`` - (DEPRECATED) Last Subject Sent::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP.
+
Use this parameter to display information only about nodes whose last
```

```
AutoSupport message had the "Subject:" line you specify.

``[-last-time <MM/DD/YYYY HH:MM:SS>]`` - (DEPRECATED) Last Time Sent::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP.
+
Use this parameter to display information only about nodes whose last
AutoSupport message was sent at the date and time you specify. Specify the
date and time in the format "MM/DD/YYYY HH:MM:SS".

``[-max-http-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum HTTPS Size::
Use this parameter to display information only about nodes that limit the
maximum size of HTTP transfers to the file size you specify.

``[-max-smtp-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum SMTP Size::
Use this parameter to display information only about nodes that limit the
maximum size of SMTP (e-mail) transfers to the file size you specify.

``[-remove-private-data {true|false}]`` - Remove Sensitive Data::
Use this parameter with the value "true" to display information only about
nodes that remove sensitive data from AutoSupport messages. Use this
parameter with the value "false" to display information only about nodes
that do not remove sensitive data.

``[-validate-digital-certificate {true|false}]`` - Validate Digital
Certificate Received [.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value "true" to display information only about
nodes that validate digital certificates they receive. Use this parameter
with the value "false" to display information only about nodes that do not
validate digital certificates.

``[-ondemand-state {enable|disable}]`` - AutoSupport OnDemand State
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about nodes that have the
AutoSupport OnDemand state you specify.

``[-ondemand-remote-diagnostics-state {enable|disable}]`` - AutoSupport
OnDemand Remote Diagnostics State [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about nodes that have the
AutoSupport OnDemand Remote Diagnostics state you specify.

``[-ondemand-server-url <text>]`` - AutoSupport OnDemand Server URL::
Use this parameter to display information only about nodes that have the
AutoSupport OnDemand Server URL you specify.
```

```
``[-compress-spool-dir {true|false}]`` - Compress AutoSupport Spool Dir
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about nodes with the
AutoSupport spool directory compression state you specify.

``[-node-specific-from {enable|disable}]`` - Node-Specific From Address
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information that matches the specified value
- enabled or disabled.

``[-validate-smtp-certificate {true|false}]`` - Validate Certificate of
SMTP Server [.advaced-priv]#(privilege: advanced)#::
Use this parameter with the value "true" to display information only about
nodes that validate certificates they receive from the SMTP server. Use
this parameter with the value "false" to display information only about
nodes that do not validate certificates from the SMTP server.

``[-smtp-encryption {none|start_tls}]`` - SMTP Encryption::
Use this parameter to display information only about nodes that use the
encryption protocol you specify to send SMTP AutoSupport messages.

== Examples

The following example displays the AutoSupport configuration for a node
named node3:
```

cluster1::> system node autosupport show -node node3
Node: node3
State: enable
SMTP Mail Hosts: smtp.example.com:25
From Address: alerts@example.com
List of To Addresses: support@example.com
List of Noteto Addresses: -
List of Partner Addresses: support@partner.com
Send AutoSupport Messages to Vendor Support: enable
Protocol to Contact Support: https
Support Proxy URL: support.proxy.example.com
Hostname Subject: true
Performance Data Enable: true
Retry Interval: 4m
Retry Count: 15
Reminder Enable: true
The Transmission Window: 1h
Last Subject Sent: WEEKLY
Last Time Sent: 3/1/2019 06:00:03
Maximum HTTP Size: 50MB
Maximum SMTP Size: 5MB
Remove Sensitive Data: false
Continue Local Collection while Disabled: true

SMTP Encryption: none

```
[[ID1781836b66cbe9c8eb50acf2e46053b9]]
= system node autosupport check show-details
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display detailed status of AutoSupport subsystem

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport check show-details` command displays the
detailed status of the AutoSupport subsystem. This includes verifying
connectivity to your vendor's AutoSupport destinations by sending test
messages and providing a list of possible errors in your AutoSupport
configuration settings.

== 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::
Selects the check results that match this parameter value. This parameter
specifies the node whose status is being displayed.

``[-check-type <Type of AutoSupport Check>]`` - AutoSupport Check Type::
Selects the check results that match this parameter value. This parameter
specifies the type of AutoSupport check being performed.

``[-status {ok|warning|failed|not-run}]`` - Status of the Check::
Selects the check results that match this parameter value. This parameter
specifies the result of this AutoSupport check.
```

```
``[-error-detail <text>]`` - Detailed Description of Error::
Selects the check results that match this parameter value. This parameter
specifies the detailed error message for this AutoSupport check.

``[-corrective-action <text>]`` - Corrective Action::
Selects the check results that match this parameter value. This parameter
specifies a description of how to correct any errors seen as part of this
AutoSupport Check

== Examples

The following example displays the detailed status of the AutoSupport
subsystem for a node named node2:
```

cluster1::> system node autosupport check show-details -node node2

```

```

Node: node2

Category: https
Component: https-put-destination
Status: ok
Detail: Successfully connected to "support.netapp.com/put".
Component: https-post-destination
Status: ok
Detail: Successfully connected to "support.netapp.com/post".

Category: smtp
Component: mail-server
Status: ok
Detail: Successfully connected to "mailhost.netapp.com".
Component: mail-server
Status: ok
Detail: Successfully connected to "sendmail.domain.com".
Component: mail-server
Status: ok
Detail: Successfully connected to "qmail.domain.com".

Category: on-demand
Component: ondemand-server
Status: ok
Detail: Successfully connected to "support.netapp.com/aods".

Category: configuration
Component: configuration
Status: ok
Detail: No configuration issues found.

```

```
[[ID2958344cce5fefc03bd984773f701b56]]
```
```

```
= system node autosupport check show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display overall status of AutoSupport subsystem

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport status check show` command displays the
overall status of the AutoSupport subsystem.

== 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::
Selects the nodes that match this parameter value. This parameter
specifies the node whose status is being displayed.

``[-http-status {ok|warning|failed|not-run}]`` - Overall Status of
AutoSupport HTTPS Destinations::
Selects the nodes that match this parameter value. This parameter
specifies whether connectivity to the AutoSupport HTTPS destination was
established.

``[-aod-status {ok|warning|failed|not-run}]`` - Overall Status of
AutoSupport OnDemand Server::
Selects the nodes that match this parameter value. This parameter
specifies the detailed description of the connectivity status to the
AutoSupport OnDemand Server.

``[-smtp-status {ok|warning|failed|not-run}]`` - Overall Status of
AutoSupport SMTP Destinations::
Selects the nodes that match this parameter value. This parameter
```

```
specifies whether connectivity to the AutoSupport mailhost was
established.

``[-config-status {ok|warning|failed|not-run}]`` - Overall Status of
AutoSupport Configuration::
Selects the nodes that match this parameter value. This parameter
specifies whether the AutoSupport configuration check succeeded or not.

``[-warning-text <text>]`` - Conditional Warning Message::
Selects the nodes that match this parameter value. This parameter
specifies how to get more details regarding the status of the AutoSupport
subsystem, in case of any errors.

== Examples

The following example displays the overall status of the AutoSupport
subsystem on a node named node2:
```

cluster1::> system node autosupport check show -node node2

```

```

On Demand
Server SMTP Configuration
---------- ---------- ----------
ok ok ok

```
[[IDcb05f078452274c0c26fba6e6ee31096]]
= system node autosupport destinations show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a summary of the current AutoSupport destinations

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport destinations show` command displays a list of
all message destinations used by AutoSupport. The command provides you
with a quick summary of all addresses and URLs that receive AutoSupport
messages from all nodes in the cluster.

== 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>|local}]`` - Node::
Use this parameter to display only destinations that receive AutoSupport
messages from the node you specify.

``[-destinations <text>,...]`` - Destinations::
Use this parameter to display only destination lists for nodes that send
AutoSupport messages to the destinations you specify.

== Examples

This example displays all of the destinations in use by the current
cluster. Each node uses the same destination for HTTPS POST, HTTPS PUT,
and e-mail notifications.
```

```
cluster1::> system node autosupport destinations show
Node
Destinations
-------------------------------------------------------------------------------
node1
https://asuppost.example.com/cgi-bin/asup.cgi
https://asupput.example.com/cgi-bin/asup.cgi
support@example.com

node2
https://asuppost.example.com/cgi-bin/asup.cgi
https://asupput.example.com/cgi-bin/asup.cgi
support@example.com
```

```
[[ID7b272383a5880de0fd9ea2178560d4fb]]
= system node autosupport history cancel
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cancel an AutoSupport Transmission.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport history cancel` command cancels an active
AutoSupport transmission. This command is used to pause or abandon a long
running delivery of an AutoSupport message. The cancelled AutoSupport
message remains available for retransmission using the
xref:{relative_path}system-node-autosupport-history-retransmit.html[system
node autosupport history retransmit] command.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node on which to cancel the AutoSupport
message. The default setting is localhost.

``-seq-num <Sequence Number>`` - AutoSupport Sequence Number::
Use this parameter to specify the sequence number of the AutoSupport
message you want to cancel.

``[-destination {smtp|http|noteto|retransmit}]`` - Destination for This
AutoSupport::
Use this parameter to specify the destination type for the AutoSupport
message you want to cancel.

== Examples

Use this command to cancel the AutoSupport message delivery with seq-num
10 to all destinations.
```

cluster1::> system node autosupport history cancel -node local -seq-num 10

```
Use this command to cancel the AutoSupport message delivery with seq-num
10 via HTTP only.
```

cluster1::> system node autosupport history cancel -node local -seq-num 10 -destination http

```
== Related Links

* xref:{relative_path}system-node-autosupport-history-
retransmit.html[system node autosupport history retransmit]




[[IDe72f23df40f0f5bef17409e7544729c8]]
= system node autosupport history retransmit
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Selectively retransmit a previously collected AutoSupport.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport history retransmit` command retransmits a
locally stored AutoSupport message.

Support personnel might ask you to run this command to retransmit an
AutoSupport message. You might also retransmit an AutoSupport message if
you run the xref:{relative_path}system-node-autosupport-history-
show.html[system node autosupport history show] command and notice that a
message was not delivered.

If you retransmit an AutoSupport message, and if support already received
that message, the support system will not create a duplicate case. If, on
the other hand, support did not receive that message, then the AutoSupport
system will analyze the message and create a case, if necessary.

Use the xref:{relative_path}system-node-autosupport-history-
show.html[system node autosupport history show] command to display the 50
most recent AutoSupport messages, which are available for retransmission.
```

```
== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node from which the AutoSupport message
is sent.

``-seq-num <Sequence Number>`` - AutoSupport Sequence Number::
Use this parameter to specify the sequence number of the AutoSupport
message to retransmit.

``-uri <text>`` - Destination to Send this AutoSupport::
Use this parameter to specify the HTTP, HTTPS, FILE, or MAILTO uniform
resource indicator (URI) to which the AutoSupport message is sent.

``[-size-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Transmit Size Limit for
this AutoSupport.::
Use this parameter to specify a size limit for the retransmitted
AutoSupport message. If the message information exceeds this limit, it
will be trimmed to fit the limit you specify. Omit the size limit or set
it to 0 to disable it, which is useful to retransmit an AutoSupport
message that was truncated by a mail or Web server due to the default size
limits.

== Examples

The following example retransmits the AutoSupport message with sequence
number 45 on the node "node1" to a support address by e-mail.
```

cluster1::> system node autosupport history retransmit -node node1 -seq-num 45 -uri mailto:support@example.com

```
== Related Links

* xref:{relative_path}system-node-autosupport-history-show.html[system
node autosupport history show]




[[IDb231b3a108b202e97f2c79a97a1e1f3d]]
= system node autosupport history show-upload-details
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display upload details of recent AutoSupport messages

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport history show-upload-details` command displays
upload details of the 50 most recent AutoSupport messages sent by nodes in
the cluster. By default, it displays the following information:

* AutoSupport Sequence Number
* Destination
* Compressed Size
* Percentage Complete
* Rate of upload
* Time Remaining



== 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  ]`` }::
Use this parameter to display the following detailed information about all
entries:
+
  * AutoSupport Sequence Number
  * Destination
  * Compressed Size
  * Percentage Complete
  * Rate of Upload
  * Time Remaining



``[-node {<nodename>|local}]`` - Node::
Use this parameter to display details of AutoSupport messages sent from
the node you specify.

``[-seq-num <Sequence Number>]`` - AutoSupport Sequence Number::
```

```
Use this parameter to display details of AutoSupport messages with the
sequence number you specify. Sequence numbers are unique to a node. Use
this parameter with the `-node` parameter to display information about an
individual message.

``[-destination {smtp|http|noteto|retransmit}]`` - Destination for this
AutoSupport::
+
NOTE: ``_http_`` destination indicates the HTTPS protocol was used to
deliver the AutoSupport.
+
Use this parameter to display details of AutoSupport messages that were
sent to the destination type you specify.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Autosupport Compressed Size::
Use this parameter to display details of AutoSupport messages of the
compressed size you specify.

``[-percent-complete <integer>]`` - Percent Complete::
Use this parameter to display the percentage completed for any active
(incomplete) AutoSupport message.

``[-upload-rate {<integer>[Bps|KBps|MBps|GBps]|unlimited}]`` - Rate of
Upload::
Use this parameter to display the rate in bytes per second that upload is
using currently, otherwise zero when not active.

``[-time-remaining <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Time remaining for Upload::
Use this parameter to display the estimated time for the transmission of
the AutoSupport message to complete.

== Examples

The following example shows the first three results output by the history
show-upload-details command. Note that "q" was pressed at the prompt.
```

cluster1::> system node autosupport history show-upload-details -node node1
Seq Percent Time
Node Num Destination Size Complete Rate Remaining
---------- ----- ---------- --------- --------- ---------- --------
node1
13
smtp 755.9KB 100 142.88KBps 0s
http 755.8KB 80 125.97KBps 10s
noteto - - - -
12

smtp - - - -
http 316.4KB 100 158.22KBps 0s
noteto 201B 100 201Bps 0s
11
smtp - - - -
http 626.1MB 100 649.56KBps 0s
noteto - - - -
Press <space> to page down, <return>> for next line, or 'q' to quit… q
9 entries were displayed.

```
[[ID8a3f2d2976a85e5b37e8ec979c1bc522]]
= system node autosupport history show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display recent AutoSupport messages

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport history show` command displays information
about the 50 most recent AutoSupport messages sent by nodes in the
cluster. By default, it displays the following information:

* AutoSupport sequence number
* Destination type, such as smtp
* Status of delivery, such as sent-successful
* Attempt count
* Time of last update



== 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.

| ``[-delivery  ]``::
Use this parameter to display destination information about each
AutoSupport message.
```

```
| ``[-detail  ]``::
Use this parameter to display trigger and subject information about each
AutoSupport message.

| ``[-instance  ]`` }::
Use this parameter to display the following detailed information about all
entries:
+
  * Trigger event
  * Subject of the message
  * Delivery URI
  * Last error
  * Compressed Size
  * Decompressed Size
  * Total Collection Time (in ms)



``[-node {<nodename>|local}]`` - Node::
Use this parameter to display only AutoSupport messages sent from the node
you specify.

``[-seq-num <Sequence Number>]`` - AutoSupport Sequence Number::
Use this parameter to display only AutoSupport messages with the sequence
number you specify. Sequence numbers are unique to a node. Use this
parameter with the `-node` parameter to display information about an
individual message.

``[-destination {smtp|http|noteto|retransmit}]`` - Destination for This
AutoSupport::
+
NOTE: ``_http_`` destination indicates the HTTPS protocol was used to
deliver the AutoSupport.
+
Use this parameter to display only AutoSupport messages that were sent to
the destination type you specify.

``[-trigger <Message Name>]`` - Trigger Event::
Use this parameter to display only AutoSupport messages that match the
trigger event you specify.

``[-last-update <MM/DD/YYYY HH:MM:SS>]`` - Time of Last Update::
Use this parameter to display only AutoSupport messages that were updated
most recently at the time you specify. Specify time in "MM/DD/YYYY
HH:MM:SS" format.
```

```
``[-status <AutoSupport general status>]`` - Status of Delivery::
Use this parameter to display only AutoSupport messages with the status
you specify. Possible statuses are:
+
  * initializing - The AutoSupport message request is being processed.
  * collection-failed - The AutoSupport collection failed. View the 'Last
Error' field of this message for more information.
  * collection-in-progress - The AutoSupport collection is in progress.
  * queued - The AutoSupport message is queued for delivery.
  * transmitting - The AutoSupport message transmission is in progress.
  * sent-successful - The AutoSupport message was sent successfully.
  * ignore - The AutoSupport message was processed successfully, but the
trigger event is not configured for delivery to the current destination
type.
  * re-queued - The AutoSupport message transmission failed, has been re-
queued, and will be retried.
  * transmission-failed - The AutoSupport message transmission failed, and
the retry limit was exceeded.
  * ondemand-ignore - The AutoSupport message was processed successfully,
but the AutoSupport On Demand server chose to ignore it.



``[-attempt-count <integer>]`` - Delivery Attempts::
Use this parameter to display only AutoSupport messages that the system
has attempted to send the number of times you specify. This parameter is
most useful when given a range, such as ">5".

``[-subject <text>]`` - AutoSupport Subject::
Use this parameter to display only AutoSupport messages of the type you
specify.

``[-uri <text>]`` - Delivery URI::
Use this parameter to display only AutoSupport messages sent to the
destination URI you specify.

``[-error <text>]`` - Last Error::
Use this parameter to display only AutoSupport messages that failed with
the "Last Error" description you specify.

``[-generated-on <MM/DD/YYYY HH:MM:SS>]`` - Time of Generation::
Use this parameter to display only AutoSupport messages that were
generated (collected) at a particular time.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - AutoSupport Compressed Size::
Use this parameter to display only AutoSupport messages of the compressed
```

```
size you specify.

``[-percent-complete <integer>]`` - Percent Complete::
Use this parameter to display the percentage completed for any active
(incomplete) AutoSupport message.

``[-upload-rate {<integer>[Bps|KBps|MBps|GBps]|unlimited}]`` - Rate of
Upload::
Use this parameter to display the rate in bytes per second that upload is
using currently, otherwise zero when not active.

``[-time-remaining <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Time Remaining for Upload::
Use this parameter to display the estimated time for the transmission of
the AutoSupport message to complete.

``[-decompressed-size {<integer>[KB|MB|GB|TB|PB]}]`` - AutoSupport
Decompressed Size::
Use this parameter to display only AutoSupport messages of the
decompressed size you specify.

``[-total-time <integer>]`` - Total Collection Time (ms)::
Use this parameter to display only AutoSupport messages of total
collection time you specify. A value is only shown when the collection has
completed.

== Examples

The following example shows the first three results output by the history
command. Note that "q" was pressed at the prompt.
```

```
cluster1::> system node autosupport history show -node node1
Seq Attempt Last
Node Num Destination Status Count Update
------------ ----- ----------- -------------------- -------- -----------------
node1 56
smtp ignore 1 11/18/2010 01:10:01
http re-queued 2 11/18/2010 02:50:07
noteto transmitting 1 11/18/2010 01:10:01
55
smtp ignore 1 11/18/2010 00:53:59
http sent-successful 3 11/18/2010 01:50:03
noteto sent-successful 1 11/18/2010 00:53:59
54
smtp ignore 1 11/17/2010 12:18:58
http sent-successful 4 11/17/2010 16:07:22
noteto sent-successful 1 11/17/2010 12:18:58
Press <space> to page down, <return>> for next line, or 'q' to quit… q
```

9 entries were displayed.

```
[[ID02a92c5a6f1522e1e815e489e239f130]]
= system node autosupport manifest show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display AutoSupport content manifest

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport manifest show` command reports what
information is contained in AutoSupport messages. The name and size of
each file collected for the message is reported, along with any errors.

== 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.

| ``[-content  ]``::
Use this parameter to display detailed information about the content of
the files contained in the report.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Use this parameter to display information about only AutoSupport messages
sent from the node you specify.

``[-seq-num <Sequence Number>]`` - AutoSupport Sequence Number::
Use this parameter to display information about only AutoSupport message
content with the sequence number you specify. Sequence numbers are unique
to a node. Use this parameter with the `-node` parameter to display
information about an individual message.
```

```
``[-prio-num <integer>]`` - Priority Order of Collection::
Use this parameter to display information about only AutoSupport message
content with the collection priority you specify. Content is collected in
order, by priority number.

``[-subsys <subsys1,subsys2,...>]`` - Subsystem::
Use this parameter to display information about only AutoSupport message
content collected by the AutoSupport subsystem you specify.

``[-cmd-tgt <Execution domain of AutoSupport content>]`` - Execution
Domain for Command::
Use this parameter to display information about only AutoSupport message
content produced in the execution domain you specify.

``[-body-file <text>]`` - The AutoSupport Content Filename for this Data::
Use this parameter to display information about only AutoSupport message
content stored in the body file with the file name you specify.

``[-cmd <text>]`` - Actual Data Being Collected::
Use this parameter to display information about only AutoSupport message
content produced by the D-Blade command, BSD command, file, or XML table
you specify.

``[-query <text>]`` - Table Query for XML Collection::
Use this parameter to display information about only AutoSupport message
content produced by the table query you specify.

``[-size-collected {<integer>[KB|MB|GB|TB|PB]}]`` - Number of Bytes
Collected::
Use this parameter to display information about only AutoSupport message
content collected in files with the file size you specify.

``[-time-collected <integer>]`` - Collection Time for this Data Item
(ms)::
Use this parameter to display information about only AutoSupport message
content collected in the amount of time you specify, in milliseconds.

``[-status <AutoSupport manifest collection status>]`` - Status of this
Data Item::
Use this parameter to display information about only AutoSupport message
content with the collection status you specify. Possible statuses are:
+
  * `requested` - The AutoSupport request has been added to the queue and
is waiting processing by the collector.
  * `working` - The AutoSupport collector is actively gathering the needed
data.
```

```
  * `file-not-found` - AutoSupport data collection failed because a
necessary file is missing.
  * `no-such-table` - The AutoSupport collector was unable to find the
requested SMF table.
  * `collection-truncated-size-limit` - AutoSupport data was truncated due
to size limits, but partial data is available.
  * `collection-truncated-file-size-limit` - AutoSupport data for a
particular data item or file was truncated due to file size limits, but
partial data is available.
  * `collection-skipped-size-limit` - AutoSupport data was skipped due to
size limits, and no data is available.
  * `collection-truncated-time-limit` - AutoSupport data was truncated due
to time limits, but partial data is available.
  * `collection-skipped-time-limit` - AutoSupport data was skipped due to
time limits, and no data is available.
  * `delivery-skipped-size-limit` - AutoSupport data was skipped at
delivery time due to size limits.
  * `collection-truncated-age` - AutoSupport data for the particular data
item or file was truncated due to age, but partial data is available.
  * `general-error` - AutoSupport data collection failed. Additional
information (if any) is in the Error String field.
  * `completed` - AutoSupport data collection is complete, and the
AutoSupport message is ready for delivery.
  * `content-not-collected-precheck` - AutoSupport content was not
collected due to pre-check function violation.
  * `content-not-collected-privacy` - AutoSupport content was not
collected because the operation is disabled in privacy mode.
  * `content-empty` - AutoSupport content was collected successfully, but
the output was empty.
  * `collection-aborted` - AutoSupport data collection was aborted.



``[-error <text>]`` - Textual Description of Error::
Use this parameter to display information about only AutoSupport message
content with the error text you specify. If data collection has failed,
the error text contains a description of the failure. If data collection
completes `successfully` , this field is empty.

``[-content-type <Type of AutoSupport content>]`` - AutoSupport Content
Type for this Data::
Use this parameter to display information about only AutoSupport message
content of the type you specify. Types supported are:
+
  * basic - Configuration data about this subsystem
  * troubleshooting - Detailed diagnostic data about this subsystem
```

```
``[-orig-size-collected {<integer>[KB|MB|GB|TB|PB]}]`` - Initial Number of
Bytes Collected::
Use this parameter to display information about only AutoSupport message
content collected in files with the original file size you specify.

``[-size-compressed {<integer>[KB|MB|GB|TB|PB]}]`` - Compressed Size::
Use this parameter to display information about only AutoSupport message
content collected in files with the compressed file size you specify.

== Examples

This example displays the content of AutoSupport message number 372 on the
node "node1".
```

cluster1::> system node autosupport manifest show -node node1 -seq-num 372
AutoSupport Collected
Node Sequence Body Filename Size Status Error
---------- --------- ---------- ---------- --------- ----------
node1 372
SYSCONFIG-A.txt 1.73KB completed
OPTIONS.txt 29.44KB completed
software_image.xml 7.56KB completed
CLUSTER-INFO.xml 3.64KB completed
autosupport.xml 12.29KB completed
autosupport_budget.xml
7.01KB completed
autosupport_history.xml
46.52KB completed
X-HEADER-DATA.TXT 717.00B completed
SYSTEM-SERIAL-NUMBER.TXT
35.00B completed
cluster_licenses.xml
3.29KB completed
cm_hourly_stats.gz 151.4KB completed
boottimes.xml 56.86KB completed
rdb_txn_latency_stats_hrly.xml
39.31KB completed
rdb_voting_latency_stats_hrly.xml
3.43KB completed
14 entries were displayed.

```
This example shows how you can use parameters to limit output to specific
fields of a specific AutoSupport message. This is helpful when
troubleshooting.
```

cluster1::> system node autosupport manifest show -node node5 -seq-num 842 -fields body-file,status,size-collected,time-collected,cmd,cmd-tgt,subsys

node seq-num prio-num subsys cmd-tgt body-file cmd size-collected time-collected status
---------- ------- -------- --------- ------- --------------- --------------- --------------- --------------- ---------
node5 842 0 mandatory dblade SYSCONFIG-A.txt "sysconfig -a" 16.44KB 256 completed
node5 842 1 mandatory dblade OPTIONS.txt options 29.67KB 3542 completed
node5 842 2 mandatory smf_table software_image.xml software_image 8.68KB 33 completed
node5 842 3 mandatory smf_table CLUSTER-INFO.xml asup_cluster_info 4.75KB 7 completed
node5 842 4 mandatory smf_table autosupport.xml autosupport 12.32KB 10 completed
node5 842 5 mandatory smf_table autosupport_budget.xml autosupport_budget 7.03KB 29 completed
node5 842 6 mandatory smf_table autosupport_history.xml autosupport_history 62.77KB 329 completed
node5 842 7 mandatory custom_fx X-HEADER-DATA.TXT "Custom function" 720.00B 3 completed
node5 842 8 mandatory custom_fx SYSTEM-SERIAL-NUMBER.TXT "Custom function" 31.00B 2 completed
node5 842 9 mandatory smf_table cluster_licenses.xml cluster_licenses 5.62KB 9 completed
node5 842 10 log_files custom_fx log_files.xml "Custom function" 13.07KB 4 completed
node5 842 11 log_files custom_fx EMS-LOG-FILE.gz "Custom function" 25.33KB 24 completed
node5 842 12 log_files dblade_file EMS-LOG-FILE-PARTNER.gz /etc/log/ems - - content-not-collected-precheck
node5 842 13 log_files dblade_file MESSAGES.gz /etc/log/messages 35.40KB 42 completed
node5 842 14 log_files dblade_file MESSAGES-PARTNER.gz /etc/log/messages - - content-not-collected-precheck
14 entries were displayed.

```
[[ID01ee17dff600072f83debf59894bbc65]]
= system node autosupport trigger modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify AutoSupport trigger configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Use the `system node autosupport trigger modify` command to enable and
disable AutoSupport messages for individual triggers, and to specify
additional subsystem reports to include if an individual trigger sends an
AutoSupport message.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node whose AutoSupport trigger
configuration is modified.

``-autosupport-message <Autosupport Message>`` - EMS Message::
Use this parameter to specify the AutoSupport trigger to modify.
```

```
AutoSupport triggers are EMS messages whose names begin with "callhome.".
However, for the purposes of this command, "callhome." is implied, does
not need to be entered, and will not be displayed in command output.

``[-to {enabled|disabled}]`` - Deliver to AutoSupport -to Addresses::
Use this parameter with the value "enabled" to enable sending AutoSupport
messages to the configured "to" addresses.

``[-noteto {enabled|disabled}]`` - (DEPRECATED) Deliver to AutoSupport
-noteto Addresses::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP. Instead, use the xref:{relative_path}event-notification-
destination-create.html[event notification destination create] command to
create an email destination and the xref:{relative_path}event-
notification-create.html[event notification create] command to map the
`important-events` system filter to the newly created event destination.
This will send an email notification for all events with a severity of
ERROR, ALERT, or EMERGENCY. For events with other severity values, a new
event filter needs to be created using the xref:{relative_path}event-
filter-create.html[event filter create] command. Filter rules can be added
to the filter using the xref:{relative_path}event-filter-rule-
add.html[event filter rule add] command. This newly created filter has to
be mapped to the event destination.
+
Use this parameter with the value "enabled" to enable sending short notes
to the configured "noteto" addresses.

``[-basic-additional <subsys1,subsys2,...>,...]`` - Additional Subsystems
Reporting Basic Info::
Use this parameter to include ``_basic_`` content from the additional
subsystems you specify. Content is collected from these subsystems in
addition to the default list of subsystems.

``[-troubleshooting-additional <subsys1,subsys2,...>,...]`` - Additional
Subsystems Reporting Troubleshooting Info::
Use this parameter to include ``_troubleshooting_`` content from the
additional subsystems you specify. Content is collected from these
subsystems in addition to the default list of subsystems.

``[-suppress {true|false}]`` - Suppress all occurrences of this trigger::
Use this parameter with the value "true" to suppress the collection when
the AutoSupport message is triggered.

== Examples
```

```
The following example enables messages to the configured "to" addresses
from the `battery.low` trigger on the node `node1` .
```

cluster1::> system node autosupport trigger modify -node node1 -autosupport-message battery.low -to enabled

```
== Related Links

* xref:{relative_path}event-notification-destination-create.html[event
notification destination create]
* xref:{relative_path}event-notification-create.html[event notification
create]
* xref:{relative_path}event-filter-create.html[event filter create]
* xref:{relative_path}event-filter-rule-add.html[event filter rule add]




[[IDcb09a08fdb0ccc822d4b695516e3d82b]]
= system node autosupport trigger show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display AutoSupport trigger configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node autosupport trigger show` command displays what system
events trigger AutoSupport messages. When a trigger event occurs, the node
may send an AutoSupport message to a predefined destination, and a short
note to another destination. The full AutoSupport message contains detail
for troubleshooting. The short message is meant for short pager or SMS
text messages.

Use the xref:{relative_path}system-node-autosupport-destinations-
show.html[system node autosupport destinations show] command to view
available destinations.

== 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.

| ``[-basic  ]``::
Use this parameter to display which subsystem information is included as
``_basic_`` content when the AutoSupport message is triggered.

| ``[-troubleshooting  ]``::
Use this parameter to display which subsystem information is included as
``_troubleshooting_`` content when the AutoSupport message is triggered.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Use this parameter to display AutoSupport triggers only on the node you
specify.

``[-autosupport-message <Autosupport Message>]`` - EMS Message::
Use this parameter to display only AutoSupport triggers with the name you
specify. AutoSupport triggers are EMS messages whose names begin with
"callhome.". However, for the purposes of this command, "callhome." is
implied, does not need to be entered, and will not be displayed in command
output.

``[-to {enabled|disabled}]`` - Deliver to AutoSupport -to Addresses::
Use this parameter with the value "enabled" to display only AutoSupport
messages that send full messages to the "to" address when triggered. Use
this parameter with the value "disabled" to display only AutoSupport
messages that do not send full messages.

``[-noteto {enabled|disabled}]`` - (DEPRECATED) Deliver to AutoSupport
-noteto Addresses::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP.
+
Use this parameter with the value "enabled" to display only AutoSupport
messages that send short notes to the "noteto" address when triggered. Use
this parameter with the value "disabled" to display only AutoSupport
messages that do not send short notes.

``[-basic-default <subsys1,subsys2,...>,...]`` - Default Subsystems
Reporting Basic Info::
```

```
Use this parameter to display only AutoSupport triggers that include in
their messages, by default, ``_basic_`` content from the subsystems you
specify.

``[-troubleshooting-default <subsys1,subsys2,...>,...]`` - Default
Subsystems Reporting Troubleshooting Info::
Use this parameter to display only AutoSupport triggers that include in
their messages, by default, ``_troubleshooting_`` content from the
subsystems you specify.

``[-additional-content <Type of AutoSupport content>,...]`` - Additional
Content Flag::
Use this parameter to display only AutoSupport triggers that have been
configured to include additional ``_basic_`` or ``_troubleshooting_``
content.

``[-basic-additional <subsys1,subsys2,...>,...]`` - Additional Subsystems
Reporting Basic Info::
Use this parameter to display only AutoSupport triggers that have been
configured to include additional ``_basic_`` content from the subsystems
you specify.

``[-troubleshooting-additional <subsys1,subsys2,...>,...]`` - Additional
Subsystems Reporting Troubleshooting Info::
Use this parameter to display only AutoSupport triggers that have been
configured to include additional ``_troubleshooting_`` content from the
subsystems you specify.

``[-suppress {true|false}]`` - Suppress all occurrences of this trigger::
Use this parameter with the value "true" to display only AutoSupport
messages that have been suppressed.

== Examples

This example shows the first page of output from the command. Note that
"q" was pressed at the prompt to quit.
```

cluster1::> system node autosupport trigger show
AutoSupport Additional
Node Message To Note To Content
--- --- --- --- ---
node1 aggr.offline enabled enabled -
node1 aggr.restricted disabled enabled -
node1 aggr.wafliron disabled enabled -
node1 bad.ram disabled disabled -
node1 battery.failure enabled enabled -
node1 battery.low disabled disabled -

node1 battery.notice enabled enabled -
node1 battery.overchg enabled enabled -
node1 battery.overtemp enabled enabled -
node1 battery.warning enabled enabled -
node1 bmc.bus disabled disabled -
node1 bmc.hb.stop disabled disabled -
node1 bmc.post disabled disabled -
node1 bootfs.chkdsk enabled enabled -
node1 c.fan enabled enabled -
node1 c.fan.fru.degraded disabled disabled -
node1 c.fan.fru.fault disabled enabled -
node1 c.fan.fru.rm disabled enabled -
node1 c.fan.fru.shut enabled enabled -
node1 ch.ps.degraded disabled disabled -
Press <space> to page down, <return> for next line, or 'q' to quit… q
20 entries were displayed.

```
== Related Links

* xref:{relative_path}system-node-autosupport-destinations-
show.html[system node autosupport destinations show]




[[ID05e860571f2f7db58cee621c8c71faca]]
= system node coredump delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all coredumps owned by a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump delete-all` command deletes either all unsaved
core dumps or all saved core files on a node. You can specify whether
saved core files or unsaved core dumps are deleted by using the optional
`-saved` parameter. If the command is issued while a core dump is being
saved, the command prompts you before stopping the save operation and
deleting the core dump.

== Parameters
```

```
``-node <nodename>`` - Node That Owns the Coredump::
This specifies the node from which core files or core dumps are to be
deleted.

``[-type {unsaved-kernel|saved-kernel|kernel|application|all}]`` - Type of
Core to delete::
This parameter specifies the type of core file to be deleted. If the type
is unsaved, all unsaved core dumps will be deleted. If the type is saved,
all saved core files will be deleted. If the type is kernel, all kernel
core files and kernel core dumps will be deleted. If the type is
application, all application core files will be deleted. If the type is
all, all core files will be deleted. The default setting is to delete only
unsaved kernel core dumps and core files.

== Examples

The following example deletes all unsaved kernel core dumps on a node
named node0:
```

cluster1::> system node coredump delete-all -node node0

```
[[IDb42656959e80f094e885f5f4260bbbde]]
= system node coredump delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a coredump

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump delete` command deletes a specified core dump.
If the command is issued while the specified core dump is being saved, the
command prompts you before stopping the save operation and deleting the
core dump.

== Parameters

``-node {<nodename>|local}`` - Node That Owns the Coredump::
This specifies the node from which core files are to be deleted.

``[-type {kernel|ancillary-kernel-segment|application}]`` - Coredump
Type::
This specifies the type of core file to be deleted. If the type is kernel,
the specified kernel core file will be deleted. If the type is
application, the specified application core file will be deleted.

``-corename <text>`` - Coredump Name::
This specifies the core file that is to be deleted.

== Examples

The following example deletes a core dump named core.101268397.2010-05-
30.19_37_31.nz from a node named node0:
```

cluster1::> system node coredump delete -node node0 -corename core.101268397.2010-05-30.19_37_31.nz

```
[[IDb2eca4b0f08d5b2fb1a5e65992c1b942]]
= system node coredump save-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Save all unsaved kernel coredumps owned by a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump save-all` saves all unsaved core dumps on a
specified node. If the node has already attempted to save the core dump by
the value set by the -save-attempts parameter, the command prompts you
before continuing. The save-attempts parameter is set by invoking the
command system node coredump config modify.

== Parameters

``-node <nodename>`` - Node That Owns the Coredump::
This specifies the node on which unsaved core dumps are to be saved.

== Examples

The following example saves all unsaved core dumps on a node named node0:
```

cluster1::> system node coredump save-all -node node0

```
[[ID0e8766839d9d69f21a4dc3ef58155823]]
= system node coredump save
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Save an unsaved kernel coredump

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump save` command saves a specified core dump. If
the node has already attempted to save the core dump by the value
specified by the `-save-attempts` parameter, the command prompts you
before continuing. The `-save-attempts` parameter is set by invoking the
command xref:{relative_path}system-node-coredump-config-modify.html[system
node coredump config modify] . A saved core dump can be uploaded to a
remote site for support analysis; see the xref:{relative_path}system-node-
coredump-upload.html[system node coredump upload] command man page for
more information.

== Parameters

``-node {<nodename>|local}`` - Node That Owns the Coredump::
This specifies the node on which the core dump is located.

``-corename <text>`` - Coredump Name::
This specifies the core dump that is to be saved.

== Examples

The following example saves a core dump named core.101268397.2010-05-
30.19_37_31.nz on a node named node0:
```

cluster1::> system node coredump save -node node0 -corename core.101268397.2010-05-30.19_37_31.nz

```
== Related Links

* xref:{relative_path}system-node-coredump-config-modify.html[system node
coredump config modify]
* xref:{relative_path}system-node-coredump-upload.html[system node
```

```
coredump upload]




[[ID4ee922fe918187700eadb1817e89cac8]]
= system node coredump show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of coredumps

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump show` command displays basic information about
core dumps, such as the core dump name, time of panic that triggered the
core dump and whether the core file is saved. You can specify optional
parameters to display information that matches only those parameters. For
example, to display the list of kernel core files, run the command with
-type kernel.

== 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.

| ``[-system  ]``::
If you specify this parameter, the command displays the following
information:
+
  * Node name
  * Core dump name
  * Core dump ID
  * Node that panicked and generated the core
  * System ID of the node that panicked and generated the core
  * Version of the core
```

```
| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node That Owns the Coredump::
If you specify both this parameter and the `-corename` parameter, the
command displays detailed information about the specified core. If you
specify this parameter by itself, the command displays information about
the core files on the specified node.

``[-type {kernel|ancillary-kernel-segment|application}]`` - Coredump
Type::
This parameter specifies the type of core files to be displayed. If the
type is kernel and the system supports segmented core files, the command
displays information about primary kernel core segment files. If the type
is kernel and the system does not support segmented core files, the
command displays information about full core files. If the type is
ancillary-kernel-segment, the command displays information about ancillary
kernel core segment files. If the type is application, the command
displays information about application core files. If no type is
specified, the command displays information about core files of type
kernel or application.

``[-corename <text>]`` - Coredump Name::
If you specify both this parameter and the `-node` parameter, the command
displays detailed information about the specified core. If you specify
this parameter by itself, the command displays information about the core
files that match the specified name.

``[-panic-node <text>]`` - Node That Generated Core::
If you specify this parameter with a node name, the command displays
information only about the core files that were generated when the
specified node panicked.

``[-panic-systemid <integer>]`` - System ID of Node That Generated Core::
If you specify this parameter, the command displays information only about
the core files that were generated when the node with the specified system
ID panicked.

``[-version <text>]`` - Data ONTAP Version of Core::
If you specify this parameter, the command displays information only about
the core files that match the specified version.

``[-panic-time <MM/DD/YYYY HH:MM:SS>]`` - Time of Panic That Generated
Core::
If you specify this parameter, the command displays information only about
```

```
the core files that were generated by a panic at the specified time.
Specify time in the format of ``_MM/DD/YYYY HH:MM:SS [+- HH:MM]_`` . You
can use ``_[+- HH:MM]_`` to specify the time range within which all core
files triggered by a panic are displayed. ``_[+- HH:MM]_`` is relative to
UTC.

``[-panic-string <text>]`` - Panic String::
If you specify this parameter, the command displays information only about
the core files that match the specified panic string.

``[-is-saved {true|false}]`` - Saved Core::
If you specify this parameter, the command displays information only about
the core dumps that are or are not saved yet to a core file.

``[-is-partial {true|false}]`` - Partial Core::
If you specify this parameter, the command displays information only about
the core dumps that are or are not partially saved.

``[-save-attempts <integer>]`` - Number of Attempts to Save Core::
If you specify this parameter, the command displays information only about
the core dumps that have the specified number of successful or failed save
attempts.

``[-space-needed {<integer>[KB|MB|GB|TB|PB]}]`` - Space Needed To Save
Core::
If you specify this parameter, the command displays information only about
the core dumps that need the specified amount of disk space to save into a
core file.

``[-size <text>]`` - Size of Core (bytes)::
If you specify this parameter, the command displays information only about
the saved core files that are of the specified size.

``[-md5-data-chksum <text>]`` - MD5 Checksum of the Compressed Data of
Core::
If you specify this parameter, the command displays information only about
the saved core files that have the specified MD5 checksum for compressed
data of the core.

``[-ancillary-segment-directory <text>]`` - Directory Holding Ancillary
Kernel Core Segments::
If you specify this parameter, the command displays information only about
the saved core files that have the specified ancillary segment directory.

== Examples

The following examples display information about the core files:
```

```
cluster1::> system node coredump show
Node Core Name Saved Panic Time
-------- ------------------------------------------ ------- -----------------
node0
core.101182345.2010-02-01.14_19_08.nz false 2/1/2010 09:19:08
Partial Core: false
Number of Attempts to Save Core: 2
Space Needed To Save Core: 4.45GB
node1
core.101268397.2010-05-30.19_37_31.nz true 5/30/2010 15:37:31
node2
core.101270930.2010-09-06.18_40_03.nz true 9/6/2010 14:40:03
node3
core.101271326.2010-09-06.19_06_18.nz true 9/6/2010 15:06:18
core.101271326.2010-09-06.19_09_49.nz true 9/6/2010 15:09:49
4 entries were displayed.
```

```
cluster1::> system node coredump show -panic-time 9/6/2010 15:00:00+3:00
Node Core Name Saved Panic Time
-------- ------------------------------------------ ------- -----------------
node3
core.101271326.2010-09-06.19_06_18.nz true 9/6/2010 15:06:18
core.101271326.2010-09-06.19_09_49.nz true 9/6/2010 15:09:49
2 entries were displayed.
```

```
[[IDda0907febd8157dfa25b574457874a95]]
= system node coredump status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display kernel coredump status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump status` command displays status information
about core dumps. The command output depends on the parameters specified
with the command. If a core dump is in the process of being saved into a
core file, the command also displays its name, the total number of blocks
that are to be saved, and the current number of blocks that are already
saved.
```

```
You can specify additional parameters to display only information that
matches those parameters. For example, to display coredump status
information about the local node, run the command with the parameter `-
node local` .

== 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.

| ``[-disks  ]``::
If you specify this parameter, the command displays the following
information:
+
  * Node name
  * Total number of disks
  * Number of spare disks
  * Number of disks used
  * Number of disks with partial cores



| ``[-spraycore  ]``::
If you specify this parameter, the command displays the following
information:
+
  * Node name
  * Whether spray cores are supported
  * Number of spray-core disks
  * Number of spray-core blocks
  * Number of disks needed for spray core
  * Estimated number of blocks needed for spray core



| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
If you specify this parameter, the command displays the following
information:
+
  * Node name
  * State of the core-dump process
```

```
  * Space available on the internal file system
  * Name of the core being saved, if applicable
  * Total number of blocks in the core being saved, if applicable
  * Number of blocks currently saved, if applicable
  * Type of core dump
  * Number of unsaved complete cores on the node
  * Number of unsaved partial cores on the node
  * Whether spray cores are supported on the node
  * Whether any spare disks are available on the node
  * Number of disks that have cores
  * Number of unsaved cores
  * Number of disks that have partial cores
  * Number of partial cores
  * Number of unused spray-core disks
  * Number of spray-core blocks
  * Number of disks available for core dumps
  * Estimated number of blocks needed for spray core
  * Number of disks needed for spray core



``[-state <text>]`` - State::
If you specify this parameter, the command displays information only about
the nodes that are in the specified core dump state. Possible values
include: nocore, idle, init, saving, and waitdump.

``[-space-available {<integer>[KB|MB|GB|TB|PB]}]`` - Space Available On
Internal Filesystem::
If you specify this parameter, the command displays information only about
the nodes that have the specified amount of available space, in bytes, on
their internal file systems.

``[-corename <text>]`` - Name of Core Being Saved::
If you specify this parameter, the command displays information only about
the node that is currently saving the specified core file name.

``[-total-blocks <integer>]`` - Total Number of Blocks in Core Being
Saved::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of blocks in the core dump being
saved.

``[-blocks-saved <integer>]`` - Number of Blocks saved::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of blocks saved.
```

```
``[-type <text>]`` - Type of Core Dump::
If you specify this parameter, the command displays information only about
the nodes that have the specified core dump type. Possible values include
zipped, sprayed, and spare.

``[-spraycore-supported {true|false}]`` - Spray Core Supported on Node::
If you specify this parameter, the command displays information only about
the nodes that do or do not support the spray method of dumping core.

``[-spares-available {true|false}]`` - Spare Disk(s) Available on Node::
If you specify this parameter, the command displays information only about
the nodes that do or do not have spare disks available.

``[-disks-used <integer>]`` - Number of Disks with Cores::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of disks that contain core dumps.

``[-unsaved-cores <integer>]`` - Number of Unsaved Complete Cores::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of complete core dumps that are
not yet saved into a core file.

``[-partial-disks <integer>]`` - Number of Disks with Partial Cores::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of disks with partial core dumps.

``[-partial-cores <integer>]`` - Number of Unsaved Partial Cores::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of partial core dumps that are
not yet saved into a core file.

``[-spraycore-disks <integer>]`` - Number of Unused Spray Core Disks::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of unused spray-core disks.

``[-spraycore-blocks <integer>]`` - Number of Spray Core Blocks::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of spray-core blocks.

``[-numdisks <integer>]`` - Total Number of Disks Available for Core
Dump::
If you specify this parameter, the command displays information only about
the nodes that have the specified total number of disks available for core
dump.

``[-blocks-needed <integer>]`` - Estimated Number of Blocks Needed for
Spray Core::
```

```
If you specify this parameter, the command displays information only about
the nodes that have the specified number of estimated blocks needed for
the spray method of dumping core.

``[-disks-needed <integer>]`` - Number of Disks Needed for Spray Core::
If you specify this parameter, the command displays information only about
the nodes that have the specified number of disks needed for the spray
method of dumping core.

``[-space-needed {<integer>[KB|MB|GB|TB|PB]}]`` - Space Needed to Save All
Unsaved Cores::
If you specify this parameter, the command displays information only about
the nodes that require the specified amount of disk space to save all
unsaved core dumps.

``[-min-free {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Free Bytes on Root
Filesystem::
If you specify this parameter, the command displays information only about
the nodes that need to have the specified number of bytes available on the
root filesystem after a core dump is saved.

== Examples

The following example displays core dump information about the node named
node0:
```

cluster1::> system node coredump status -node node0 -instance
Node: node0
State: idle
Space Available On Internal Filesystem: 132.1GB
Name of Core Being Saved: -
Total Number of Blocks in Core Being Saved: -
Number of Blocks saved: -
Type of core dump: spray
Number of Unsaved Complete Cores: 0
Number of Unsaved Partial Cores: 1
Space Needed To Save All Unsaved Cores: 4.81GB
Minimum Free Bytes On Internal Filesystem: 250MB

```
[[IDc9f922da4e00c0a821cb8a7427cd80c3]]
= system node coredump trigger
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Make the node dump system core and reset

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command triggers a Non-maskable Interrupt (NMI) on the specified node
via the Service Processor of that node, causing a dirty shutdown of the
node. This operation forces a dump of the kernel core when halting the
node. LIF migration or storage takeover occurs as normal in a dirty
shutdown. This command is different from the -dump parameter of the
`system node shutdown` , xref:{relative_path}system-node-halt.html[system
node halt] , or xref:{relative_path}system-node-reboot.html[system node
reboot] command in that this command uses a control flow through the
Service Processor of the remote node, whereas the -dump parameter uses a
communication channel between ONTAP running on the nodes. This command is
helpful in cases where ONTAP on the remote node is hung or does not
respond for some reason. If the panic node reboots back up, then the
generated coredump can be seen by using the xref:{relative_path}system-
node-coredump-show.html[system node coredump show] command. This command
works for a single node only and the full name of the node must be entered
exactly.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node for which you want to trigger a
coredump.

== Examples

The following example triggers a NMI via the Service Processor and causes
node2 to panic and generate a coredump. Once node2 reboots back up, the
command xref:{relative_path}system-node-coredump-show.html[system node
coredump show] can be used to display the generated coredump.
```

```
cluster1::> set advanced

Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y

cluster1::*> system node coredump trigger -node node2

Warning: The Service Processor is about to perform an operation that will cause
a dirty shutdown of node "node2". This operation can
cause data loss. Before using this command, ensure that the cluster
will have enough remaining nodes to stay in quorum. To reboot or halt
a node gracefully, use the "system node reboot" or "system node halt"
command instead. Do you want to continue? {yes|no}: yes

Warning: This operation will reboot the current node. You will lose this login
session. Do you want to continue? {y|n}: y

cluster1::*>
cluster1::> system coredump show
Node:Type Core Name Saved Panic Time
--------- ---------------------------------------- ----- -----------------
node2:kernel
core.1786429481.2013-10-04.11_18_37.nz false 10/4/2013 11:18:37
Partial Core: false
Number of Attempts to Save Core: 0
Space Needed To Save Core: 3.60GB
1 entries were displayed.

cluster1::>
```

```
== Related Links

* xref:{relative_path}system-node-halt.html[system node halt]
* xref:{relative_path}system-node-reboot.html[system node reboot]
* xref:{relative_path}system-node-coredump-show.html[system node coredump
show]




[[IDda60616d9b5de3133bc58c207dae9edc]]
= system node coredump upload
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Upload a coredump to a remote site
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and might be removed in a future release
of ONTAP. Use "link:system-node-autosupport-invoke-core-upload.html[system
node autosupport invoke-core-upload] " instead.



The `system node coredump upload` command uploads a saved core file to a
specified URL. You should use this command only at the direction of
technical support.

== Parameters

``-node {<nodename>|local}`` - Node That Owns the Coredump::
This specifies the node on which the core file is located.

``[-type {kernel|ancillary-kernel-segment|application}]`` - Coredump
Type::
This specifies the type of core files to be uploaded. If the type is
kernel, kernel core files will be uploaded. If the type is application,
application core file will be uploaded.

``-corename <text>`` - Coredump Name::
This specifies the name of the core file that is to be uploaded.

``[-location <text>]`` - URL for Coredump Upload Directory::
This specifies the URL to which the core file is to be uploaded. If this
parameter is not specified, the command uploads the core file to the
location specified by the `-upload-location` parameter of the
xref:{relative_path}system-node-coredump-config-modify.html[system node
coredump config modify] command. The following protocols are supported:
ftp and http. (By default, the location is set to ftp://ftp.netapp.com/to-
ntap/)

``[-casenum <integer>]`` - Case Number::
This specifies the support case number that will be prefixed to the core
file name at the destination. The case number is critical information for
quick and automated processing of the received core file.

== Examples
```

```
The following example uploads a core file named core.07142005145732.2010-
10-05.19_03_41.nz on a node named node0 to the default location. The
support case number is 2001234567.
```

cluster1::> system node coredump upload -node node0 -corename core.07142005145732.2010-10-05.19_03_41.nz -casenum 2001234567

```
== Related Links

* xref:{relative_path}system-node-autosupport-invoke-core-
upload.html[system node autosupport invoke-core-upload]
* xref:{relative_path}system-node-coredump-config-modify.html[system node
coredump config modify]




[[ID0602704f5d500f1187277305edc6505b]]
= system node coredump config modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify coredump configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump config modify` command modifies the cluster's
core dump configuration.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node whose coredump configuration you want to
modify.

``[-sparsecore-enabled {true|false}]`` - Enable Sparse Cores::
If you set this parameter to true, the command enables sparse cores. A
sparse core omits all memory buffers that contain only user data.

``[-min-free {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Free Bytes On Root
```

```
Filesystem::
If you specify this parameter, the command displays the number of bytes
that need to be made available in the root file system after saving the
core dump. If the minimum number of bytes cannot be guaranteed, core dumps
are not generated. The default setting is 250 MB.

``[-coredump-attempts <integer>]`` - Maximum Number Of Attempts to Dump
Core::
If you specify this parameter, the command displays the maximum number of
times the system will attempt to generate a core dump when encountering
repeated disk failures. The default setting is 2.

``[-save-attempts <integer>]`` - Maximum Number Attempts to Save Core::
If you specify this parameter, the command displays the maximum number of
times the system will attempt to save a core dump.The default setting is
2.

``[-save-onstartup {true|false}]`` - Enable Auto Save of Coredumps on
Startup::
If you set this parameter to true, the system will automatically start
saving the core dump after reboot.

``[-upload-location <text>]`` - URL for Coredump Upload Directory::
+
NOTE: This option is deprecated and might be removed in a future release
of ONTAP. Use the `-uri` parameter of the "link:system-node-autosupport-
invoke-core-upload.html[system node autosupport invoke-core-upload] "
command instead.
+
If you specify this parameter, the system uploads the core dumps to the
specified URL. The following protocols are supported: ftp and http. (The
default setting is ftp://ftp.netapp.com/to-ntap/.)

== Examples

The following example sets the maximum number of core dump attempts to 5
and the maximum number of save attempts to 5:
```

cluster1::> system node coredump config modify -coredump-attempts 5 -save-attempts 5

```
== Related Links

* xref:{relative_path}system-node-autosupport-invoke-core-
upload.html[system node autosupport invoke-core-upload]
```

```
[[ID0ae1e615508028c9830d611f3f81121e]]
= system node coredump config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display coredump configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump config show` command displays basic information
about a cluster's core dump configuration, such as whether sparse cores
are enabled, minimum number of free bytes on the root volume file system
that need to be available after saving the core files, maximum number of
times the process attempts to generate a core dump when encountering
repeated disk failures, maximum number of times the process attempts to
save a core dump, the URL to which core dumps are uploaded, and whether
core dumps are automatically saved when a node restarts.




== 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>|local}]`` - Node::
If you specify this parameter, the command displays the coredump
configuration information of the specified node.

``[-sparsecore-enabled {true|false}]`` - Enable Sparse Cores::
If you specify this parameter, the command displays only the coredump
information that matches the specified spare core setting. A sparse core
```

```
omits all memory buffers that contain only user data.

``[-min-free {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Free Bytes On Root
Filesystem::
If you specify this parameter, the command displays only the core dump
information that matches the specified number of bytes that need to be
made available in the root file system after saving the core dump.

``[-coredump-attempts <integer>]`` - Maximum Number Of Attempts to Dump
Core::
If you specify this parameter, the command displays only the core dump
information that matches the specified maximum number of times the system
will attempt to generate a core dump when encountering repeated disk
failures.

``[-save-attempts <integer>]`` - Maximum Number Attempts to Save Core::
If you specify this parameter, the command displays only the coredump
information that matches the maximum number of times the system will
attempt to save a core dump.

``[-save-onstartup {true|false}]`` - Enable Auto Save of Coredumps on
Startup::
If you specify this parameter, the command displays only the coredump
information that matches the specified configuration of whether the system
will automatically start saving the core dump after reboot.

``[-upload-location <text>]`` - URL for Coredump Upload Directory::
+
NOTE: This option is deprecated and might be removed in a future release
of ONTAP. Use the `-uri` parameter of the "link:system-node-autosupport-
invoke-core-upload.html[system node autosupport invoke-core-upload] "
command instead.
+
If you specify this parameter, the command displays only the core dump
information that matches the specified URL where core dumps are uploaded.

== Examples

The following example displays information about the cluster's core dump
configuration:
```

cluster1::> system node coredump config show
Sparse Min Max Max On
Core Free Dump Save Startup
Node Enabled Bytes Attempts Attempts Enabled Coredump Location
----- ------- -------- -------- -------- ------- ----------
node0

true 250MB 2 2 true ftp://ftp.example.com/to-example/
node1
true 250MB 2 2 true ftp://ftp.example.com/to-example/
node2
true 250MB 2 2 true ftp://ftp.example.com/to-example/
node3
true 250MB 2 2 true ftp://ftp.example.com/to-example/
4 entries were displayed.

```
== Related Links

* xref:{relative_path}system-node-autosupport-invoke-core-
upload.html[system node autosupport invoke-core-upload]




[[IDbd6fbc162a6ba8004db0149714a8f6f9]]
= system node coredump external-device save
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Save a core dump to an external USB device

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node coredump external-device save` command saves a specified
core dump to an external USB device plugged into the port specified by the
`-device` parameter.

External USB device requirements:

* A device formatted with the FAT32 filesystem.
* A device formatted with the ext2 filesystem with the largefile flag set.
+ This can be done using the following command on a Linux host: mkfs.ext2
-T largefile <device_name>
* The command xref:{relative_path}system-node-coredump-show.html[system
node coredump show] can be used to determine the size of the core dump.
```

```
== Parameters

``-node {<nodename>|local}`` - Node That Owns the Coredump [.advaced-
priv]#(privilege: advanced)#::
This specifies the node on which the core dump is located.

``-device {usb0|usb1}`` - Device [.advaced-priv]#(privilege: advanced)#::
This specifies which external USB device connected to the node to save the
core dump, for example: usb0. Currently, only usb0 is supported.  + usb0
refers to the first external USB device connected to the node (regardless
of port).

``-corename <text>`` - Coredump Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the core dump that is to be saved.

== Examples

The following example saves a core dump named core.101268397.2010-05-
30.19_37_31.nz on node1 to external USB device usb0:
```

cluster1::> system node coredump external-device save -node node1 -device usb0 -corename core.101268397.2010-05-30.19_37_31.nz

```
== Related Links

* xref:{relative_path}system-node-coredump-show.html[system node coredump
show]




[[IDed945fcaf37be4411654df93e70f72b9]]
= system node coredump external-device show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of files on an external USB device

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description
```

```
The `system node coredump external-device show` command displays basic
information about files on an external USB device, such as the filename
and size.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``-node {<nodename>|local}`` - Node That Owns the Coredump [.advaced-
priv]#(privilege: advanced)#::
This parameter selects the node that has files that are to be displayed on
the external USB device.

``[-device {usb0|usb1}]`` - Device [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the external USB device, for example:
usb0. Currently, only usb0 is supported.  + usb0 refers to the first
external USB device connected to the node (regardless of port).

``[-corename <text>]`` - Coredump Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the core dump file for which the information is
displayed.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of Core [.advaced-
priv]#(privilege: advanced)#::
If specified, the command displays information only about the core files
that are of the specified size.

== Examples

The following example displays information about core files stored on
external USB devices:
```

```
cluster1::> system node coredump external-device show
Node Device Coredump Name
----------------------- ------- -------------------------------------------
node1 usb0 core.537051938.2017-10-26.23_52_15.nz
Size: 13074695581 bytes (12.18) GB
```

```
[[ID9f66948652c8f3ae7f5b23e600d086c]]
= system node coredump reports delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an application core report

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node coredump reports delete` command deletes the specified
application core report.

== Parameters

``-node {<nodename>|local}`` - Node That Owns the Coredump::
This specifies the node from which reports are to be deleted.

``-reportname <text>`` - Report Name::
This specifies the report that is to be deleted.

== Examples

The following example shows how a report named notifyd.1894.80335005.2011-
03-25.09_59_43.ucore.report is deleted from a node named node0:
```

cluster1::> system node coredump reports delete -node node0 -reportname notifyd.1894.80335005.2011-03-25.09_59_43.ucore.report

```
[[ID10f5fc378504c5a17f73fb0d0a926b70]]
= system node coredump reports show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of application core reports

*Availability:* This command is available to _cluster_ administrators at
```

```
the _admin_ privilege level.

== Description

The `system node coredump reports show` command displays basic information
about application core reports, such as the report name and time of the
panic that triggered the application core dump. You can specify optional
parameters to display information that matches only those parameters. For
example, to display the list of reports in the local node, run the command
with -node local.

== 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>|local}]`` - Node That Owns the Coredump::
Selects information about all the reports on the specified node. If you
specify both this parameter and the -reportname parameter, the command
displays detailed information about the specified report.

``[-reportname <text>]`` - Report Name::
Selects information about the reports that match the specified name. If
you specify both this parameter and the -node parameter, the command
displays detailed information about the specified report.

``[-panic-node <text>]`` - Node That Generated Core::
Selects information about the reports that were generated by the specified
node.

``[-panic-systemid <integer>]`` - System ID of Node That Generated Core::
Selects information about the reports that were generated by the node with
the specified system ID.

``[-version <text>]`` - Data ONTAP Version of Core::
Selects information about the reports that match the specified version.

``[-panic-time <MM/DD/YYYY HH:MM:SS>]`` - Time of Panic That Generated
Core::
Selects information about the reports that were generated by a panic at
the specified time. Specify time in the format of ``_MM/DD/YYYY HH:MM:SS
```

```
[+- HH:MM]_`` . You can use ``_[+- HH:MM]_`` to specify the time range
within which all core files triggered by a panic are displayed. ``_[+-
HH:MM]_`` is relative to UTC.

``[-panic-string <text>]`` - Panic String::
Selects information about the reports that match the specified panic
string.

== Examples

The following example displays information about the reports:
```

cluster1::> system node coredump reports show
Node Report Name Panic Time
-------- ---------- ----------
node0 notifyd.1894.80335005.2011-03-25.09_59_43.ucore.report 3/25/2011 09:59:43

```
[[ID8d70bd3bf55feac5457d69e9078a605f]]
= system node coredump reports upload
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Upload an application core report to a remote site

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and might be removed in a future release
of ONTAP. See core report information in the SmartSoft tool.



The `system node coredump reports upload` command uploads an application
report to a specified URL. You should use this command only at the
direction of technical support.

== Parameters

``-node {<nodename>|local}`` - Node That Owns the Coredump::
This specifies the node on which the report is located.
```

```
``-reportname <text>`` - Report Name::
This specifies the name of the report that is to be uploaded.

``[-location <text>]`` - URL for Coredump Upload Directory::
This specifies the URL to which the report is to be uploaded. The
following protocols are supported: ftp and http. (By default, the location
is set to ftp://ftp.netapp.com/to-ntap/)

``[-casenum <integer>]`` - Case Number::
This specifies the support case number that is be prefixed to the core
file name at the destination. The case number is critical information for
quick and automated processing of the received core file.

== Examples

The following example shows how a report named notifyd.1894.80335005.2011-
03-25.09_59_43.ucore.bz2 is uploaded on a node named node0 to the default
location. The support case number is 2001234567.
```

cluster1::> system node coredump reports upload -node node0 -corename notifyd.1894.80335005.2011-03-25.09_59_43.ucore.bz2 -casenum 2001234567

```
[[ID77901999457b063234051247c01f9a1f]]
= system node coredump segment delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all core segments on a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command deletes all the core segments on a node.

== Parameters

``-node {<nodename>|local}`` - Node::
This specifies the node on which to delete the core segments.

== Examples

This deletes all the core segments for ``_node1_`` .
```

cluster1::> system node coredump segment delete-all -node node1

```
[[IDa58f1eb24ec105d4f6ff23bb781c64b5]]
= system node coredump segment delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a core segment

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command deletes a core segment.

== Parameters

``-node {<nodename>|local}`` - Node::
This specifies the node on which to delete the core segments.

``-segment <text>`` - Core Segment::
This specifies the core segment to delete. The pathname is relative to the
coredump directory. If a directory is specified, all core segment files
within it are deleted. If the directory is empty, it is deleted.

``[-owner-node <text>]`` - Node That Owns the Core Segment File::
This specifies the node that owns the core segment. Use this parameter
only in takeover mode to delete a partner's coredump segment.

== Examples

This deletes all core segments in the directory, ``_core.151708240.2012-
01-11.05_56_52_`` .
```

cluster1::> system node coredump segment delete -node node1 -segment core.151708240.2012-01-11.05_56_52

```
[[ID41068fe8ea09d84ab9acbdc25fd16007]]
= system node coredump segment show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display a list of core segments

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the following information about core segments:

* name of the core segment directory
* time of the panic that generated the core segment
* total number of core segment files
* core segment file name



== 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  ]`` }::
Displays the following details:
+
  * Core segment file name
  * Node that owns the core segment file
  * System ID of the node that generated the core
  * MD5 checksum of the compressed data of the core segment file
  * Name of the core segment
  * Total number of core segments for the core file
  * Timestamp of the panic that triggered the core segment



``[-node {<nodename>|local}]`` - Node::
Selects information about the core segments on the specified node.

``[-segment <text>]`` - Core Segment::
Selects information about the specified core segment. If segment is a
directory, the command displays the information for the first core segment
file. If segment is a file, the command displays the file information.

``[-owner-node <text>]`` - Node That Owns the Core Segment File::
```

```
Selects information about the core segments owned by the specified node.
This parameter should only be used in takeover mode to display information
about the partner's core segments.

``[-panic-system-id <integer>]`` - System ID of Node That Generated Core::
Selects information about the core segments that were generated when the
node with the specified system ID panicked.

``[-md5-data-chksum <text>]`` - Md5 Checksum of the Compressed Data of the
Core Segment::
Selects information about the core segments whose data segment's MD5
checksum matches the specified checksum.

``[-segment-name <text>]`` - Name of the Core Segment::
Selects information about the core segments with the specified name.

``[-total-segment-count <integer>]`` - Number of Segments Generated::
Selects information about the core segments with the specified name.

``[-panic-time <MM/DD/YYYY HH:MM:SS>]`` - Time of Panic That Generated
Core::
Selects information about the core segments that were generated by a panic
at the specified time.

``[-size <text>]`` - Size of Core Segment (bytes)::
Selects information about the core segments that are of the specified
size.

``[-panic-string <text>]`` - Panic String of Panic That Generated Core::
Selects information about the core segments that match the specified panic
string.

== Examples

The example below displays the core segments on ``_node1_`` .
```

cluster1::> system node coredump segment show -node node1
Node: node1
Segment Directory: core.118049106.2012-01-05.17_11_11
Panic Time: 1/5/2012 12:11:11
Number of Segments: 2
Segment File Name:
core.118049106.2012-01-05.17_11_11.nvram.nz
core.118049106.2012-01-05.17_11_11.ontap.nz
2 entries were displayed.

```
The example below displays detailed information a specific core segment
file on ``_node1_`` .
```

cluster1::> system node coredump segment show -node node1 -segment core.118049106.2012-01-05.17_11_11.ontap.nz -instance
Node: node1
Core Segment: core.118049106.2012-01-05.17_11_11.ontap.nz
Node That Owns the Core Segment File: node1
System ID of Node That Generated Core: 118049106
Md5 Checksum of the Compressed Data of the Core Segment: 1a936d805dcd4fd5f1180fa6464fdee4
Name of the Core Segment: ontap
Number of Segments Generated: 2
Time of Panic That Generated Core: 1/5/2012 12:11:11

```
[[ID833b86d9dcc6f92425c35b73c1c4f28a]]
= system node environment sensors show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the sensor table

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node environment sensors show` command displays the following
information:

* Node name
* Sensor name
* Sensor state
* Sensor value
* Sensor units
* Critically Low threshold for the sensor
* Warning Low threshold for sensor
* Warning High threshold for sensor
* Critically High threshold for the sensor
* FRU name (detailed view only)



== 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>|local}]`` - Node::
Selects information about the sensors on the specified node.If this
parameter is specified with the -name parameter, the command displays
information only about the specified sensor.

``[-name <text>]`` - Sensor Name::
Selects information about the sensors that have the specified name.If this
parameter is specified with the -node parameter, the command displays
information only about the specified sensor.

``[-fru <text>]`` - FRU::
Selects information about the sensors associated with the specified Field
Replaceable Unit (FRU).

``[-type {fan|thermal|voltage|current|battery-
life|discrete|fru|nvmem|counter|minutes|percent|agent|unknown}]`` - Sensor
Type::
Selects information about the sensors that have the specified sensor type.
Possible values vary among platforms but may include ``_fan_`` ,
``_temperature_`` , ``_thermal_`` and ``_voltage_`` .

``[-units <text>]`` - Value Units::
Selects information about the sensors that have readings displayed in the
specified units of measure. Possible values vary among platforms but may
include ``_RPM_`` , ``_C_`` and ``_mV_`` .

``[-state {normal|warn-low|warn-high|crit-low|crit-
high|disabled|uninitialized|init-failed|not-
available|invalid|retry|bad|not-
present|failed|ignored|fault|unknown|nonreadable}]`` - Sensor State::
Selects information about the sensors that have the specified state.
Possible values vary among platforms but may include ``_normal_`` ,
``_warn_lo_`` , ``_warn_hi_`` , ``_crit_lo_`` , ``_crit_hi_`` and
``_failed_`` .

``[-discrete-state {normal|warn-low|warn-high|crit-low|crit-
```

```
high|disabled|uninitialized|init-failed|not-
available|invalid|retry|bad|not-
present|failed|ignored|fault|unknown|nonreadable}]`` - Discrete Sensor
State::
Selects information about the discrete-valued sensors that are in the
specified state. A discrete-valued sensor has a set of possible discrete
values rather than a range of possible values. For example, a presence
sensor which has the discrete values PRESENT and NOT_PRESENT is a
discrete-valued sensor. Possible values vary among platforms but may
include ``_normal_`` and ``_failed_`` .

``[-value <integer>]`` - Last Sensor Value::
Selects information about the sensors that have the specified sensor
value.

``[-discrete-value <text>]`` - Discrete Sensor Value::
Selects information about the discrete-valued sensors that have the
specified discrete value. Possible values vary among sensors but may
include ``_PRESENT_`` , ``_NOT_PRESENT_`` , ``_ON_`` , ``_OFF_`` ,
``_OK_`` and ``_FAULT_`` .

``[-crit-low <integer>]`` - Critical Low Threshold::
Selects information about the sensors that have the specified critically
low threshold.

``[-warn-low <integer>]`` - Warning Low Threshold::
Selects information about the sensors that have the specified warning-low
threshold.

``[-warn-hi <integer>]`` - Warning Hi Threshold::
Selects information about the sensors that have the specified warning-high
threshold.

``[-crit-hi <integer>]`` - Critical Hi Threshold::
Selects information about the sensors that have the specified critically
high threshold.

``[-inactive {true|false}]`` - Show Inactive Sensors::
Specify ``_true_`` to include inactive sensors in the output. By default,
only sensors with the value ``_false_`` are shown.

``[-hidden {true|false}]`` - Show Hidden Sensors::
Specify ``_true_`` to include hidden sensors in the output. By default,
only sensors with the value ``_false_`` are shown.

== Examples
```

```
The following example displays information about all sensors on a cluster
named cluster1:
```

```
cluster1::> system node environment sensors show
Node Sensor State Value/Units Crit-Low Warn-Low Warn-Hi Crit-Hi
---- --------------------- ------ ----------- -------- -------- ------- -------
mynode
Partner IO Pre
NOT_PRESENT
Partner Ctrl Pre
PRESENT
PSU2 Over Curr
OK
PSU2 Over Volt
OK
PSU2 Over Temp
OK
PSU2 Fault
OK
PSU2 DC OK
OK
PSU2 Input OK
OK
PSU2 ON
ON
PSU2 Fan2 Fault
OK

Node Sensor State Value/Units Crit-Low Warn-Low Warn-Hi Crit-Hi
---- --------------------- ------ ----------- -------- -------- ------- -------
mynode
PSU2 Fan2 Speed
normal 15400 RPM 3000 3500 - 25500
PSU2 Fan1 Fault
OK
PSU2 Fan1 Speed
normal 15700 RPM 3000 3500 - 25500
PSU2 Curr
normal 28000 mA - - - -
PSU2 Temp
normal 29 C 0 5 51 61
PSU2 Present
PRESENT
PSU1 Over Curr
OK
PSU1 Over Volt
OK
PSU1 Over Temp
OK

Node Sensor State Value/Units Crit-Low Warn-Low Warn-Hi Crit-Hi
---- --------------------- ------ ----------- -------- -------- ------- -------
mynode
```

PSU1 Fault
OK
PSU1 DC OK
OK
PSU1 Input OK
OK
PSU1 ON
ON
PSU1 Fan2 Fault
OK
PSU1 Fan2 Speed
normal 15600 RPM 3000 3500 - 25500
PSU1 Fan1 Fault
OK
PSU1 Fan1 Speed
normal 16200 RPM 3000 3500 - 25500
PSU1 Curr
normal 27000 mA - - - -
PSU1 Temp
normal 28 C 0 5 51 61

Node Sensor State Value/Units Crit-Low Warn-Low Warn-Hi Crit-Hi
---- --- --- --- --- --- --- ---
mynode
PSU1 Present
PRESENT
Battery 3.3V
normal 3400 mV 3025 3100 3500 3575
AUX 3.3V
normal 3328 mV 3024 3104 3504 3568
STBY 12V
normal 12152 mV 10478 10602 13392 13516
STBY 5V
normal 4979 mV 4602 4696 5310 5404
STBY 3.3V
normal 3375 mV 3025 3100 3500 3575
12V
normal 12152 mV 10478 10726 13268 13516
5V
normal 5003 mV 4602 4696 5310 5404
3.3V
normal 3375 mV 3025 3100 3500 3575

```
[[IDa7a99e4d71ff815666f39045cd230fc0]]
= system node external-cache modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify external cache settings.
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node external-cache modify` command can be used to modify the
following attributes of external-cache for a node:

* is-enabled
* is-rewarm-enabled
* is-mbuf-inserts-enabled
* pcs-size
* is-hya-enabled



== Parameters

``-node {<nodename>|local}`` - Node::
This specifies the node on which the modifications need to be made.

``[-is-enabled {true|false}]`` - Is Enabled?::
Enables external-cache module (Flash Cache Family) on the storage system.
Valid values for this option are true and false. If external-cache
hardware is present, then this option will enable external-cache
functionality in WAFL. If no hardware is present, this option will enable
external-cache pcs (Predictive Cache Statistics). The default value for
this option is false.

``[-is-rewarm-enabled {true|false}]`` - Is Rewarm On?::
Specifies whether an external-cache module should attempt to preserve data
across reboots. Valid values for this option are true and false. This
option applies only to cache hardware with persistent media. It does not
apply to Predictive Cache Statistics (PCS). Enabling this option will
marginally increase the duration of system boot and shutdown, but it will
reduce or eliminate the time required for cache warming. The default value
for this option is determined by the cache hardware type. The option is
disabled by default.

``[-is-mbuf-inserts-enabled {true|false}]`` - Is Mbuf Inserts On?::
Specifies whether the external-cache module allows insert of mbuf data as
part of a network write. In rare cases, inserting mbuf data may cause
excessive CPU usage. We provide this workaround to disable the behavior,
if necessary. Do not change the value of this option unless directed to do
so by technical support. The data from the mbuf network writes can still
be stored in the external cache, but only after a subsequent disk read of
```

```
that data.

``[-pcs-size <integer>]`` - PCS Size::
Controls the size of the cache emulated by external-cache PCS. Valid
values for this option are integers between 16 and 16383. This option is
only used when PCS is enabled. The default value for this option is chosen
automatically based on the amount of memory in the controller, and the
upper limit is further restricted on controllers with smaller amounts of
memory.

``[-is-hya-enabled {true|false}]`` - Is HyA Caching Enabled?::
Specifies whether the external-cache module allows caching of blocks
targeted for hybrid aggregates. This option is set to true by default when
the external-cache is enabled.

== Examples
```

cluster::> system node external-cache modify -node node1 -is-enabled true

```
The command enables the external-cache feature on `node1` .




[[ID1fc280e9298ec9cd6d2de9006da9333e]]
= system node external-cache show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display external cache settings.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node external-cache show` command displays external-cache
information for each of the nodes available.

== Parameters

{ ``[-fields <fieldname>,...]``::
Valid values for this option are `{node|is-enabled|is-rewarm-enabled|is-
```

```
mbuf-inserts-enabled|pcs-size|is-hya-enabled}` . Specifying the value will
display all entries that correspond to it.

| ``[-instance  ]`` }::
This option does not need an input value. Specifying this option will
display the information about all the entries.

``[-node {<nodename>|local}]`` - Node::
Specify this parameter to display external-cache parameters that match the
specified node.

``[-is-enabled {true|false}]`` - Is Enabled?::
Valid values for this option are true and false. Specifying the value will
display all entries that correspond to it.

``[-is-rewarm-enabled {true|false}]`` - Is Rewarm On?::
Valid values for this option are true and false. Specifying the value will
display all entries that correspond to it.

``[-is-mbuf-inserts-enabled {true|false}]`` - Is Mbuf Inserts On?::
Valid values for this option are true and false. Specifying the value will
display all entries that correspond to it.

``[-pcs-size <integer>]`` - PCS Size::
Valid values for this option are integers between 16 and 16383. Specifying
the value will display all entries that correspond to it.

``[-is-hya-enabled {true|false}]`` - Is HyA Caching Enabled?::
Valid values for this option are true and false. Specifying the value will
display all entries that correspond to it.

== Examples
```

cluster1::> system node external-cache show -node node1
Node: node1
Is Enabled: false
Is rewarm on: false
Is Mbuf inserts on: true
PCS size: 256
Is hya caching enabled: true

```
Displays the external-cache information about `node1` in a list format.




[[ID70fac8be0247105c31979740531ea912]]
= system node firmware download
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Download motherboard firmware and system diagnostics

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node firmware download` command downloads new system firmware
to the boot device. A reboot followed by the 'update_flash' command at the
firmware prompt is required for the firmware to take effect.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the node or nodes on which the firmware is to be updated.

``-package <text>`` - Package URL [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the URL that provides the location of the package
to be fetched. Standard URL schemes, including HTTP, HTTPS, FTP and FILE,
are accepted. The FILE URL scheme can be used to specify location of the
package to be fetched from an external device connected to the storage
controller. Currently, only USB mass storage devices are supported. The
USB device is specified as ``_\file://usb0/<filename>_`` . Typically, the
file name is ``_image.tgz_`` . The package must be present in the root
directory of the USB mass storage device. The HTTPS URL scheme requires
that you install the HTTPS server certificate on the system by using the
command "security certificate install -type server-ca".

== Examples

The following example downloads firmware to node-01 from a web server:
```

cluster1::*> system node firmware download -node node-01 -package http://example.com/serviceimage.zip

```
[[ID27edd752392f4238901ce9cd7c294011]]
= system node firmware bootmedia show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the BIOS/LOADER firmware versions stored in the boot media and
running on the system

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the BIOS and LOADER versions of system firmware
installed on the boot media, as well as the versions currently running on
the system.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-instance  ]`` }::
Displays detailed information about all the controllers in the cluster.

``[-node {<nodename>|local}]`` - Node::
Selects information about the specified controller.

``[-board-ID <text>]`` - Board Type::
Selects information about the controllers with the specified board type.

``[-modelName <text>]`` - Model Name::
Selects information about the controllers with the specified model name.

``[-bootmedia-bios-version <text>]`` - BIOS Firmware Version::
Selects information about the controllers with the specified BIOS firmware
version installed on the boot media.

``[-bootmedia-loader-version <text>]`` - LOADER Firmware Version::
Selects information about the controllers with the specified LOADER
```

```
firmware version installed on the boot media.

``[-running-bios-version <text>]`` - BIOS version::
Selects information about the controllers with the specified BIOS firmware
version currently running on the system.

``[-running-loader-version <text>]`` - Loader version::
Selects information about the controllers with the specified LOADER
firmware version currently running on the system.

== Examples

The following example displays the BIOS and LOADER versions installed on
the bootmedia as well as the versions currently running on the system:
```

cluster1::> system node firmware bootmedia show
Node Model Name BIOS Version LOADER Version
---------- ---------- ---------- ----------
node01
FAS2820
Boot Media: 19.5 8.4.0
Running: 19.4 8.3.0
node02
FAS2820
Boot Media: 19.5 8.4.0
Running: 19.4 8.3.0

```
[[ID28827989b00e63c395b7ceb4a1413ead]]
= system node hardware nvram-encryption modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Configure NVRAM device encryption

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node hardware nvram-encryption modify` command configures the
encryption feature for the NVRAM or NVMEM data that is destaged to non-
volatile flash storage.

NOTE: This feature might be restricted in some countries due to local
regulations concerning encrypted data.



== Parameters

``-node {<nodename>|local}`` - Node::
Specifies the node containing the NVRAM or NVMEM subsystem.

``[-is-enabled {true|false}]`` - Is Encryption Enabled::
Specifies whether the NVRAM or NVMEM encryption is disabled or enabled.

== Examples

The following commands enable or disable the NVRAM encryption:
```

cluster1::> system node hardware nvram-encryption modify -node node1 -is-enabled false
cluster1::> system node hardware nvram-encryption modify -node node1 -is-enabled true

```
[[ID570bcd851e7a5851151a55a8ea7d9c31]]
= system node hardware nvram-encryption show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Show NVRAM device encryption information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node hardware nvram-encryption show` command displays the
configuration of the encryption feature for the NVRAM or NVMEM data that
is destaged to non-volatile flash storage.

== 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>|local}]`` - Node::
If this parameter is specified, the command displays information about the
NVRAM encryption configuration on the specified node.

``[-nvram-device-name <text>]`` - NVRAM Device Name::
If this parameter is specified, the command displays information about the
NVRAM encryption configuration for the specified NVRAM device. Current
platforms only support one device - NVRAM.0.

``[-is-supported {true|false}]`` - Is Encryption Support::
If this parameter is specified, the command displays information about the
NVRAM encryption configuration for platforms that support it.

``[-is-enabled {true|false}]`` - Is Encryption Enabled::
If this parameter is specified, the command displays information about the
NVRAM encryption configuration for the NVRAM or NVMEM devices where the
device has the specified enabled value.

``[-key-id <text>]`` - Key ID of the Encryption Key::
If this parameter is specified, the command displays information about the
NVRAM encryption configuration with the specified encryption Key ID used
to encrypt the NVRAM or NVMEM data on flash storage.
```

```
== Examples

The following example displays information about the NVRAM encryption
configuration on all nodes of the cluster:
```

cluster1::> system node hardware nvram-encryption show
Node: node1
NVRAM-Device: NVRAM.0
Supported: true
Enabled: true
Key-ID: 0000000000000000000000000000000000000000000000000000000000000000
Node: node2
NVRAM-Device: NVRAM.0
Supported: true
Enabled: true
Key-ID: 0000000000000000000000000000000000000000000000000000000000000000
2 entries are displayed.

```
[[ID7a517669fcb9064817531ae8290904d3]]
= system node hardware tape drive show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays information about tape drives

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the following information about tape drives:

* Node to which the tape drive is attached
* Device ID of the tape drive
* Description of the tape drive
* NDMP path of the tape drive



== 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>|local}]`` - Node::
Displays detailed information about tape drives on the specified node.

``[-device-id <text>]`` - Device ID::
Selects information about the tape drive that has the specified device ID.

``[-description <text>]`` - Description::
Selects information about the tape drive or drives that have the specified
description.

``[-wwn <text>]`` - World Wide Name::
Selects information about the tape drive that has the specified worldwide
name.

``[-serial-number <text>]`` - Serial Number::
Selects information about the tape drive that has the specified serial
number.

``[-ndmp-path <text>,...]`` - NDMP Path::
Selects information about the tape drive or drives that have the specified
NDMP path.

== Examples

The following example displays information about all tape drives in the
cluster:
```

cluster1::> system node hardware tape drive show
Node Device Id Drive Description NDMP Path
------ --------- ---------- ----------
cluster1
brocade-247-198:3.126L1 nrst0l nrst0m nrst0h nrst0a
IBM LTO 4 ULTRIUM rst0l rst0m rst0h rst0a
urst0l urst0m urst0h urst0a
brocade-247-198:3.126L2 nrst1l nrst1m nrst1h nrst1a
IBM LTO 4 ULTRIUM rst1l rst1m rst1h rst1a
urst1l urst1m urst1h urst1a
brocade-247-198:3.126L3 nrst2l nrst2m nrst2h nrst2a
IBM LTO 4 ULTRIUM rst2l rst2m rst2h rst2a
urst2l urst2m urst2h urst2a
brocade-247-198:3.126L4 nrst3l nrst3m nrst3h nrst3a
IBM LTO 4 ULTRIUM rst3l rst3m rst3h rst3a

urst3l urst3m urst3h urst3a
brocade-247-198:3.126L6 nrst5l nrst5m nrst5h nrst5a
SONY SDX-400C rst5l rst5m rst5h rst5a
urst5l urst5m urst5h urst5a
5 entries were displayed.

```
[[ID7bdf67d0767e484e274f9bf09da327da]]
= system node hardware tape library show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about tape libraries

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the following information about tape libraries:

* Node to which the tape library is attached
* Device ID of the tape library
* Description of the tape library
* NDMP path of the tape library



== 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>|local}]`` - Node::
Displays detailed information about tape libraries on the specified node.

``[-device-id <text>]`` - Device ID::
Selects information about the tape library that has the specified device
ID.
```

```
``[-description <text>]`` - Description::
Selects information about the tape library or libraries that have the
specified description.

``[-wwn <text>]`` - World Wide Name::
Selects information about the tape library that has the specified
worldwide name.

``[-serial-number <text>]`` - Serial Number::
Selects information about the tape library that has the specified serial
number.

``[-ndmp-path <text>]`` - NDMP Path::
Selects information about the tape library or libraries that have the
specified NDMP path.

== Examples

The following example displays information about all tape libraries
attached to the cluster:
```

cluster1::> system node hardware tape library show
Node Device Id Drive Description NDMP Path
------ --------- -------------------- ----------------
cluster1-00
0b.125L1 HP MSL G3 mc1
Series
0c.125L1 HP MSL G3 mc0
Series
2 entries were displayed.

```
[[ID3181dd57cba7c289c45ff68c786b9e94]]
= system node hardware unified-connect modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the Fibre Channel and converged networking adapter configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `system node hardware unified-connect modify` command changes the
adapter configuration. Any changes to the adapter mode or type will
require a reboot for the changes to take effect. The adapter must also be
offline before you can make any changes.

The adapter argument is in the form Xy where X is an integer and y is a
letter. For example: ``_4a_``

For a target adapter, use the xref:{relative_path}network-fcp-adapter-
modify.html[network fcp adapter modify] command to bring the adapter
offline.

For an initiator adapter, use the `system node run local storage disable
adapter` command to take the adapter offline.

The `-mode` parameter refers to the mode of the adapter and can be either
``_fc_`` or ``_cna_`` .

The `-type` parameter refers to the FC-4 type of the adapter and can be
``_initiator_`` , ``_target_`` , or ``_fcvi_`` .

The `-force` parameter suppresses confirmation prompts.

NOTE: The adapter type ``_fcvi_`` is supported only on platforms with FCVI
adapters.



== Parameters

``-node {<nodename>|local}`` - Node::
Specifies the node of the adapter.

``-adapter <text>`` - Adapter::
Specifies the adapter.

``[-m, -mode {fc|cna}]`` - Configured Mode::
Specifies the mode.

``[-t, -type {initiator|target|fcvi}]`` - Configured Type::
Specifies the FC-4 type.

``[-f, -force <true>]`` - Force::
Suppresses warnings and confirmation prompts.

== Examples
```

cluster1::> system node hardware unified-connect modify -node node1 -adapter 0d -mode cna

```
Configures the mode of adapter 0d on node1 to CNA.

== Related Links

* xref:{relative_path}network-fcp-adapter-modify.html[network fcp adapter
modify]




[[IDf38fbe87552d9560f99d56507694a72e]]
= system node hardware unified-connect show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays information about Fibre Channel and converged networking adapters

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command manages Fibre Channel and converged networking adapters used
by the storage subsystem. Use the command to show the current mode and FC-
4 type of adapters or the capabilities of adapters.

== 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.

| ``[-capability  ]``::
If this parameter is specified, the command displays the capabilities of
the adapters.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays information about
Fibre Channel and converged networking adapters on the specified node.

``[-adapter <text>]`` - Adapter::
If this parameter is specified, the command displays information about the
specified adapter.

``[-current-mode {fc|cna}]`` - Current Mode::
If this parameter is specified, the command displays adapters configured
to the specified mode.

``[-current-type {initiator|target|fcvi}]`` - Current Type::
If this parameter is specified, the command displays adapters configured
to the specified FC-4 type.

``[-pending-mode {fc|cna}]`` - Pending Mode::
If this parameter is specified, the command displays adapters configured
to the specified mode on the next reboot.

``[-pending-type {initiator|target|fcvi}]`` - Pending Type::
If this parameter is specified, the command displays adapters configured
to the specified FC-4 on the next reboot.

``[-status-admin <text>]`` - Administrative Status::
If this parameter is specified, the command displays adapters with the
specified status.

``[-supported-modes {fc|cna}]`` - Supported Modes::
The list of modes that the adapter supports.

``[-supported-fc-types {initiator|target|fcvi}]`` - Supported FC Types::
The list of FC-4 types the adapter supports when configured into fc mode.

``[-supported-cna-types {initiator|target|fcvi}]`` - Supported CNA Types::
The list of FC-4 types the adapter supports when configured into cna mode.

== Examples

The following example displays information about all Fibre Channel and
converged networking adapters in the cluster:
```

cluster1::> system node hardware unified-connect show
Current Current Pending Pending Admin
Node Adapter Mode Type Mode Type Status
---------- ------- ------- --------- ------- --------- -----------
node1 0c fc initiator - - online

node1 0d fc initiator - - online
node1 3a fc target - - online
node1 3b fc target - - online
node1 4a cna target - - online
node1 4b cna target - - online
node1 6a fc target - - online
node1 6b fc target - - online
node2 0c fc initiator - - online
node2 0d fc initiator - - online
node2 3a fc target - - online
node2 3b fc target - - online
node2 4a cna target - - online
node2 4b cna target - - online
node2 6a cna target - - online
node2 6b cna target - - online
16 entries were displayed.

```
[[ID7b498e36be283c404542cf4ae3c7f345]]
= system node image abort-operation
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Abort software image 'update' or 'get' operation

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node image abort-operation` command aborts software
installation ("update") or download ("get") operation on the specified
node.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the node on which to abort the operation.

== Examples

The following example aborts the software installation operation on a node
named ``_node1_`` .
```

cluster1::> system node image abort-operation -node node1

```
[[IDb8724783d8ce4a41fc27f6e8af72dc1d]]
= system node image get
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Fetch a file from a URL

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command fetches a file from the specified URL and stores it in the
/mroot/etc/software directory.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node that fetches and stores the package.

``-package <text>`` - Package URL [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the URL that provides the location of the package
to be fetched. Standard URL schemes, including HTTP, HTTPS, FTP and FILE,
are accepted. The FILE URL scheme can be used to specify the location of
the package to be fetched from an external device connected to the storage
controller. Currently, only USB mass storage devices are supported. The
USB device is specified as ``_\file://usb0/<filename>_`` . Typically, the
file name is ``_image.tgz_`` . The package must be present in the root
directory of the USB mass storage device. The HTTPS URL scheme requires
that you install the HTTPS server certificate on the system by using the
command "security certificate install -type server-ca".

``[-replace-package <true>]`` - Replace the Local File [.advaced-
priv]#(privilege: advanced)#::
Specifies whether an existing package is deleted and replaced with a new
package. If you enter this command without using this parameter, its
effective value is false and an existing package is not replaced with the
new one. If you enter this parameter without a value, it is set to true
and an existing package is replaced with the new one.

``[-rename-package <text>]`` - Rename the File [.advaced-priv]#(privilege:
advanced)#::
```

```
Use this parameter to enter a package name that is different from the file
name in the URL.

``[-background <true>]`` - Run in the background [.advaced-
priv]#(privilege: advanced)#::
This parameter allows the operation to run in the background. The progress
of the operation can be checked with the command system image show-update-
progress. If this command is entered without using this parameter, its
effective value is false and the operation runs in the foreground. If this
parameter is used without a value, it is set to true.

== Examples
```

system image get http://example.com/image.tgz -rename-package image2.tgz -replace-package

```
[[IDd942c9671b2b96edebb983221ae3d61e]]
= system node image modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify software image configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node image modify` command sets the default software image on
a specified node. The default software image is the image that is run when
the node is started. A node holds two software images; when you set one as
the default image, the other image is automatically unset as the default.
Conversely, if you unset a software image as the default, the other image
is automatically set as the default.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the node on which the software image is located.

``-image {image1|image2|remote}`` - Image Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the software image that is to be set or unset as the
default.

``[-isdefault {true|false}]`` - Is Default Image [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether the specified image is the default.

== Examples

The following example sets the software image named image2 as the default
image on a node named node0.
```

node::> system node image modify -node node0 -image image2 -isdefault true
Default Image Changed.

```
[[ID039ba9eb5c2ddb4e99fdb31c9aea9814]]
= system node image show-update-progress
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show progress information for a currently running update

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node image show-update-progress` command displays the progress
of a software-image update initiated by using the
xref:{relative_path}system-node-image-update.html[system node image
update] command. The command displays progress until the update completes;
you can also interrupt it by pressing Ctrl-C.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This optionally specifies the name of a node whose image-update progress
is to be displayed.

``[-follow <true>]`` - Follow the Progress in the Foreground [.advaced-
priv]#(privilege: advanced)#::
Do use not use background processing for this command. If you do not use
this parameter, the value is ``_true_`` .

== Examples

The following example displays image-update progress:
```

node::> system node image show-update-progress

ERROR: command failed: There is no update/install in progress

```
== Related Links

* xref:{relative_path}system-node-image-update.html[system node image
update]
```

```
[[ID9a7bc54b0b6933c4c93525d3113672be]]
= system node image show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display software image information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node image show` command displays information about software
images. By default, the command displays the following information:

* Node name
* Image name
* Whether the image is the default image
* Whether the image is the current image
* Software version
* Installation date

To display detailed information about a specific software image, run the
command with the -node and -image parameters. The detailed view adds
information about the kernel image path, and the root file system image
path.

You can specify additional parameters to select specific information. For
example, to display information only about software images that are
currently running, run the command with the `-iscurrent true` parameter.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
```

```
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects information about the software images on the specified node. If
this parameter and the -image parameter are both used, the command
displays detailed information about the specified software image.

``[-image {image1|image2|remote}]`` - Image Name [.advaced-
priv]#(privilege: advanced)#::
Selects information about the software images that match the specified
name. If this parameter and the -node parameter are specified, the command
displays detailed information about the specified software image.

``[-isdefault {true|false}]`` - Is Default Image [.advaced-
priv]#(privilege: advanced)#::
Selects information about the software images with the specified default
setting.

``[-iscurrent {true|false}]`` - Is Current Image [.advaced-
priv]#(privilege: advanced)#::
Selects information about the software images that have the specified
currency value.

``[-kernel-path <text>]`` - Kernel Image Path [.advaced-priv]#(privilege:
advanced)#::
Selects information about the software images that have the specified
kernel image path.

``[-rootfs-path <text>]`` - Root Filesystem Image Path [.advaced-
priv]#(privilege: advanced)#::
Selects information about the software images that have the specified root
file system image path.

``[-version <text>]`` - Software Version [.advaced-priv]#(privilege:
advanced)#::
Selects information about the software images that have the specified root
file system image path.

``[-installdate <MM/DD/YYYY HH:MM:SS>]`` - Install Date [.advaced-
priv]#(privilege: advanced)#::
Selects information about the software image that have the specified
installation date. Specify the date in the format MM/DD/YYYY HH:MM:SS [+-
HH:MM].

== Examples
```

```
The following example displays information about the software images on a
node named node1:
```

cluster1::> system node image show -node node1
Is Is Install
Node Image Default Current Version Date
------ ------ -------- ------- --------- ------------------
node1
image1 true true 8.0 8/20/2009 17:42:42
image2 false false 8.0 6/26/2009 17:44:50
2 entries were displayed.

```
[[IDd09a4cff62611000bc1bd80e47a06160]]
= system node image update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Perform software image upgrade/downgrade

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node image update` command downloads the software image from a
specified location and updates the alternate software image (that is, the
image that is not currently running on the node). By default, validation
of the software image is not performed. Use the "-validate-only" parameter
to validate the software image first, before performing the update on the
cluster nodes.

At the advanced privilege level, you can specify whether to disable
version-compatibility checking.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the node on which the software image is located.

``-package <text>`` - Package URL [.advaced-priv]#(privilege: advanced)#::
This specifies the location from which the software image is to be
downloaded. The location can be specified in any of the following ways:
+
```

```
  * As an HTTP URL in the form ``_\http://host_name[:port]/path_to_file_``
. For instance, ``_\http://example.com/downloads/image.tgz_`` . The
management utility prompts you for a user name and password before
beginning the download.
+
NOTE: If you use HTTP to transfer software images, be aware that the
management utility does not check whether the Web server is password
protected; if it is not, press Enter at the prompt for user name and
password.
+

  * As an HTTPS URL in the form
``_\https://host_name[:port]/path_to_file_`` . For instance,
``_\https://example.com/downloads/image.tgz_`` .
+
NOTE: The HTTPS URL scheme requires that you install the HTTPS server
certificate on the system by using the command "security certificate
install -type server-ca".
+

  * As an FTP URL in the form ``_\ftp://host_name[:port]/path_to_file_`` .
For instance, ``_\ftp://example.com/downloads/image.tgz_`` . If required,
the management utility prompts you for a user name and password before
beginning the download.
  * As a filename of a package left behind by a previous installation, or
a package fetched using xref:{relative_path}system-node-image-
get.html[system node image get] . For example, ``_image.tgz_`` . Available
packages can be displayed using xref:{relative_path}system-node-image-
package-show.html[system node image package show] .
  * As a path to a package in a mounted file system in the form
``_\file://localhost/path_to_file_`` . For example,
``_\file://localhost/mroot/etc/software/image.tgz_`` .
  * The FILE URL scheme can be used to specify the location of the package
to be fetched from an external device connected to the storage controller.
Currently, only USB mass storage devices are supported. The USB device is
specified as ``_\file://usb0/<filename>_`` . Typically, the file name is
``_image.tgz_`` . The package must be present in the root directory of the
USB mass storage device.



``[-replace {image1|image2}]`` - Image to Replace [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies the image that is to be replaced when the node
is booted from the network.
```

```
``[-setdefault <true>]`` - Set Newly Updated Image as Default [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether to set the newly updated image as the
default image (that is, the image that runs the next time the node is
restarted). Note that for this parameter to work correctly, the cluster
must be in quorum when the image is updated.

``[-replace-package <true>]`` - Replace the Local File [.advaced-
priv]#(privilege: advanced)#::
Specifies whether an existing package is deleted and replaced with a new
package. If this command is entered without using this parameter, its
effective value is false and an existing package is not replaced with the
new one. If this parameter is used without a value, it is set to true and
an existing package is replaced with the new one.

``[-rename-package <text>]`` - Rename the File [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to enter a package name that is different from the file
name in the URL.

``[-background <true>]`` - Run in the Background [.advaced-
priv]#(privilege: advanced)#::
This parameter will allow the operation to run in the background. The
progress of the operation can be checked with the command
xref:{relative_path}system-node-image-show-update-progress.html[system
node image show-update-progress] . If this command is entered without
using this parameter, its effective value is false and the operation will
run in the foreground. If this parameter is used without a value, it is
set to true.

``[-validate-only <true>]`` - Validate the Package before Installation
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to validate the package. Validation consists of
verifying whether there is enough space on the system to install the
package, verifying the checksum for each component within the package and
so on. Validation usually takes from 30 to 60 minutes. If you specify this
parameter, the package will be validated only, not installed.

== Examples

The following example updates the software image on a node named node0
from a software package located at
ftp://ftp.example.com/downloads/image.tgz:
```

node::> system node image update -node node0 -package ftp://ftp.example.com/downloads/image.tgz -setdefault true

```
== Related Links

* xref:{relative_path}system-node-image-get.html[system node image get]
* xref:{relative_path}system-node-image-package-show.html[system node
image package show]
* xref:{relative_path}system-node-image-show-update-progress.html[system
node image show-update-progress]




[[IDbdb0e1ff4c55af8f7a49cad3a5db3cfa]]
= system node image package delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a software package

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The delete command will delete the specified software package.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
The package will be deleted from the repository belonging to the node
specified with this parameter. The local node is used as the default if
this parameter is omitted.

``-package <text>`` - Package File Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the package to be deleted.

== Examples
```

::> system image package delete image.tgz
1 entry was deleted.

```
[[ID000dd85d30f9a8fa315009f2ad086715]]
= system node image package show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display software package information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The package show command displays details of the software packages
residing on the storage controller.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects which node's packages are displayed. The local node is the default
if this parameter is omitted.

``[-package <text>]`` - Package File Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies which package's information will be displayed.

== Examples
```

cluster1::> system image package show
Package
Node Repository Package File Name
---- ---------- ----------
node-01
mroot

image.tgz
1 entries were displayed.

```
[[ID6cf0cbd25691b43b0278e5ee42a21064]]
= system node image package external-device delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete file on external device

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The delete command deletes the specified file on the external device.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
The file is deleted from the external device of the node specified with
this parameter. If this parameter is omitted, then the local node is used
as the default node.

``-package <text>`` - File Name [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the file to be deleted.

``-device {usb0|usb1}`` - Device [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the name of the external device. Currently, only
usb0 is supported.  + usb0 refers to the first external USB device
connected to the node (regardless of port).

== Examples
```

::> system image package external delete -package image.tgz

```
[[ID9d59ff0d09acfa14bc4b9a3c37454c45]]
= system node image package external-device show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display file listing on external device

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The external-device show command displays files residing on the external
storage device.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter selects the node that has files that are to be displayed on
the external storage device. If this parameter is omitted, then the local
node is the default node.

``[-package <text>]`` - File Name [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the file for which the information is displayed.

``[-device {usb0|usb1}]`` - Device [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the external device. Currently, only
usb0 is supported.  + usb0 refers to the first external USB device
connected to the node (regardless of port).

== Examples
```

cluster1::> system image package external-device show
Node Device Package File Name
---------- --------- ----------
node-01 usb0 image.tgz
node-01 usb0 netboot.tgz
2 entries were displayed.

```
[[ID528e1355311722162403a02fcb0617ea]]
= system node internal-switch show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display onboard switch attributes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node internal-switch show` command is used to display the
internal switch state information and the link status.

== 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>|local}]`` - Node::
Use this parameter to specify the node the switch resides on.

``[-switch-id <integer>]`` - Switch::
Use this parameter to specify the switch id. For example, 1.

``[-port-id <integer>]`` - Port::
Use this parameter to specify the port id. For example, 0.

``[-port-name <text>]`` - Port Name::
Use this parameter to specify the port name. For example, e0M.
```

```
``[-auto-admin <Auto-negotiation setting>]`` - Auto-Negotiation
Administrative::
Use this parameter to show the auto-negotiation administrative setting.
'enable' or 'disable'.

``[-auto-op <Auto-negotiation setting>]`` - Auto-Negotiation Operational::
Use this parameter to show the auto-negotiation operational setting.
'unknown', 'complete', 'incomplete', 'failded' or 'disabled'.

``[-duplex-admin <Duplex>]`` - Duplex Mode Administrative::
Use this parameter to show the duplex mode administrative setting. 'half'
or 'full'.

``[-duplex-op <Duplex>]`` - Duplex Mode Operational::
Use this parameter to show the duplex mode operational setting. 'half' or
'full'.

``[-speed-admin <Link speed>]`` - Speed Administrative::
Use this parameter to show the speed administrative setting. '10', '100'
or '1000'.

``[-speed-op <Link speed>]`` - Speed Operational::
Use this parameter to show the speed operational setting. '10', '100' or
'1000'.

``[-link <Link Status>]`` - Link State::
Use this parameter to show the link state, 'up' or 'down'.

``[-up-admin <Link Status>]`` - Up Administrative::
Use this parameter to show the up administrative setting, 'up' or 'down'.

``[-fc-op <Flow control>]`` - Flow Control Operational::
Use this parameter to show the flow control operational setting, 'full',
'send', 'receive' or 'none'.

== Examples

The example shows the attributes of the internal switch 0 on the node
Node1.
```

cluster1::> system node internal-switch show -node Node1 -switch-id 0
Auto-Negot Duplex Speed(Mbps)
Port Role Link Admin/Oper Admin/Oper Admin/Oper
--- --- --- --- --- ---
Node: Node1 , Switch: 0
0 sw-wrench up enable/complete full/full 1000/1000

1 sw-locked-wrench down enable/incomplete full/half 100/10
2 sw-e0M up enable/complete full/full 1000/1000
3 sw-e0P down enable/incomplete full/half 100/10
4 sw-midplane-1 down enable/incomplete full/half 100/10
5 sw-expander-1 up enable/unknown full/full 100/100
6 sw-sp-1 up enable/unknown full/full 100/100
7 entries were displayed.

```
[[ID6770b5461c90eb77ac1f1b3fa210805f]]
= system node internal-switch dump stat
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display onboard switch port statistics counter

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node internal-switch dump stat` command is used to display the
counter information of the internal switch ports.

== 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::
Use this parameter to specify the node the switch resides on.

``[-switch-id <integer>]`` - Switch::
Use this parameter to specify the switch id. For example, 1.

``[-port-id <integer>]`` - Port::
Use this parameter to specify the port id. For example, 0.

``[-stat-id <text>]`` - Counter Name::
```

```
Use this parameter to specify the counter name.

``[-valued <integer>]`` - Counter Value::
Use this parameter to show the value of specified counter.

``[-port-name <text>]`` - Port Name::
Use this parameter to specify the port name. For example, e0M.

== Examples

The following example shows partial counter information of the internal
switch 0 on Node1
```

cluster1::> system node internal-switch dump stat -node Node1 -switch-id 0
Port Port Name Counter Value
---- --- --- ---
Node: Node1 , Switch: 0
0 sw-wrench 1024ToMaxOctets 22480201
0 sw-wrench 128To255Octets 119552
0 sw-wrench 256To511Octets 345587
0 sw-wrench 512To1023Octets 1250437
0 sw-wrench 64Octets 803025

```
[[ID3217571bb52880a72290c7206d88e192]]
= system node nfs usage show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show NFS usage in the local node

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node nfs usage show ` command displays the NFS usage
information in the local node. The display output shows the number of RPC
calls received per protocol on the local node. Usage is collected whenever
there is any NFS traffic. These values are not persistent and will reset
when the node reboots.

== Examples

The following example displays the NFS usage information that does not
have any NFS usage.
```

::*> system node nfs usage show
Node: node1
v3: 0
v4: 0

```
The following example displays the NFS usage information with v3 usage.
```

::*> system node nfs usage show
Node: node1
v3: 5
v4: 0

```
The following example displays the NFS usage information with v4 usage.
```

::*> system node nfs usage show
Node: node1
v3: 0
v4: 14

```
[[IDcccd36a4c6bc5c7a755642b691f7c423]]
= system node power on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Power nodes on

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command switches on the power of the main controller of the specified
node. This command works for a single node only and the full name of the
node must be entered exactly.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node whose power you want to switch on.

== Examples

The following example switches on the power of node2.
```

cluster1::> set advanced

Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y

cluster1::*>
cluster1::*> system node power on -node node2

cluster1::*>

```
[[ID961add883e671276c1e200edaee2392c]]
= system node power show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the current power status of the nodes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the power status of the main controller in each node
across the cluster.

== 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>|local}]`` - Node::
This optional parameter specifies the name of a node for which information
is to be displayed. If this parameter is not specified, the command
displays information about all nodes in the cluster.

``[-status {on|off}]`` - Current Power Status::
If the -status parameter is specified, the command only lists information
about the node with the power status value you enter.

== Examples

The following example displays power status of all the nodes in cluster1.
```

cluster1::> system node power show

Node Status
--- ---

node1 on
node2 on
2 entries were displayed.

cluster1::>

```
[[IDacc37b241f1b993367e60fb3287572ba]]
= system node root-mount create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a mount from one node to another node's root volume.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node root-mount create` command produces a root-mount from one
node in the cluster to another node's root volume. The root-mount is
marked for immediate creation by a background task. Use the
xref:{relative_path}system-node-root-mount-show.html[system node root-
mount show] command to view the current status of root-mount or verify
task completion.

== Parameters

``-node <nodename>`` - Owner of the Root-mount::
The node name where the root-mount will be created.

``-root-node <nodename>`` - Root-mount Destination Node::
The node name that the root-mount will access.

== Examples

The following example shows the creation of a root-mount from
``_cluster1::nodeA_`` to ``_cluster1::nodeB_`` and the verification of the
successful completion.
```

cluster1::> system node root-mount show
This table is currently empty.

cluster1::> system node root-mount create -node nodeA -root-node nodeB

```
cluster1::> system node root-mount show
Node Root Node State Last Error
----------------- ----------------- ----------- --------------------------------
nodeA nodeB ready
```

```
== Related Links

* xref:{relative_path}system-node-root-mount-show.html[system node root-
mount show]




[[IDe29302348d7b6c6128cf906a34fa634f]]
= system node root-mount delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a mount from one node to another node's root volume.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node root-mount delete` command removes a root-mount from one
node in the cluster to another node's root volume. The root-mount is
marked for immediate deletion by a background task. Use the
xref:{relative_path}system-node-root-mount-show.html[system node root-
mount show] command to view the current status of root-mount or verify
task completion.

== Parameters

``-node <nodename>`` - Owner of the Root-mount::
The node which has the mount.

``-root-node <nodename>`` - Root-mount Destination Node::
The node accessed by the mount.

== Examples

This example shows the deletion of a root-mount from ``_cluster1::nodeA_``
to ``_cluster1::nodeB_`` and the verification of the command's successful
completion.
```

cluster1::> system node root-mount show

Node Root Node State Last Error
---------- ---------- ---------- ----------
nodeA NodeB ready

cluster1::> system node root-mount delete -node nodeA -root-node nodeB

cluster1::> system node root-mount show
This table is currently empty.

```
== Related Links

* xref:{relative_path}system-node-root-mount-show.html[system node root-
mount show]




[[ID88df752ab6c197bdfdda5e898b9cb47b]]
= system node root-mount show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the existing mounts from any node to another node's root volume.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node root-mount show` command displays the status of current
root-mounts from any node to another node's root volume. These root-mounts
are used by cluster services to access data on other nodes in the cluster.
These root-mounts are not pre-created, but are created as they are needed.
They can also be manually created or deleted.

== 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>]`` - Owner of the Root-mount::
Selects information about root-mounts that exist on the specified node.

``[-root-node <nodename>]`` - Root-mount Destination Node::
Selects information about root-mounts that connect to the specified node.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Mount Creation Time::
Selects information about root-mounts that were created at the specified
time.

``[-state <Mount State>]`` - State of the Root-Mount::
Selects information about root-mounts that have the specified state. The
states are:
+
  * ``_unknown_`` : The state of the root-mount is being determined.
  * ``_initializing_`` : A root-mount was found and needs   testing to
determine the correct state.
  * ``_mount-requested_`` : The root-mount has been requested, but is not
ready.
  * ``_mounting_`` : The root-mount is being created, but is not ready.
  * ``_ready_`` : The root-mount is ready to be used.
  * ``_not-responding_`` : The root-mount exists but is not responding.
  * ``_does-not-exist_`` : No root-mount is possible to this node's root
volume.
  * ``_ha-busy_`` : The root-mount is busy pending completion of an HA
event.
  * ``_clean-up-requested_`` : The root-mount is being deleted.
  * ``_cleaning-up_`` : The root-mount is being deleted.
  * ``_create-error_`` : The root-mount could not be created.



``[-last-error <text>]`` - Last Error::
Selects information about root-mounts that have the specified last-error
value.

== Examples

 + The following example shows the default state of the root-mounts on a
cluster that is not using root-node services:
```

cluster1::> system node root-mount show
This table is currently empty.

```
 + The following example displays the root-mounts that exist for a cluster
that has ``_nodeA_`` mounted to ``_nodeB_`` , and ``_nodeB_`` mounted to
``_nodeA_`` :
```

cluster1::> system node root-mount show
Node Root Node State Last Error
---------- ---------- ---------- ----------
nodeA nodeB ready
nodeB nodeA ready
2 entries were displayed.

```
[[IDee7280cfe94436bbdce45f8f6a14f82c]]
= system node upgrade-revert show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display upgrade/revert node status.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node upgrade-revert show` command displays information about
the status of upgrades or reversions. In case of upgrade to a patch
release, for example from "9.12.1P1" to "9.12.1P2", there are no upgrade
tasks to run, and therefore this command will not return any output. If an
upgrade has failed, this command enables you to determine which phase of
the upgrade contains the failed upgrade task and the reason for the
failure.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display status information only about upgrades or
reversions that are slated to occur on the nodes you specify.

``[-upgrade-version <integer>]`` - Cluster Upgrade Version [.advaced-
priv]#(privilege: advanced)#::
Selects status information about upgrades or reversions that are to the
version number you specify.

``[-startup-phase {pre-root|pre-apps|post-apps}]`` - Startup Phase
[.advaced-priv]#(privilege: advanced)#::
Selects status information about upgrades or reversions that are slated to
occur during the startup phase you specify. Startup phases are:
+
  * pre-root - Upgrade is applied before mroot is mounted
  * pre-apps - Upgrade is applied before other cluster apps are started
  * post-apps - Upgrade is applied after all RDB apps are online



``[-status <Upgrade/Revert Execution Status>]`` - Execution Status
[.advaced-priv]#(privilege: advanced)#::
Selects status information about upgrades or reversions that have the
execution status you specify. Execution statuses are:
+
  * prepared - Ready to upgrade
  * applied - Successful upgrade
  * reverted - Successful reversion
  * failed - Unsuccessful upgrade or reversion
  * aborted - Unsuccessful upgrade or reversion
  * skipped - Upgrade or reversion was skipped for that phase
  * locked - Upgrading or reverting



``[-status-msg <text>]`` - Status Message [.advaced-priv]#(privilege:
advanced)#::
Selects status information about upgrades or reversions that have the
status message you specify. The status message displays the current status
of the phase with which it appears.

``[-direction {upgrade|revert}]`` - Upgrade/Revert Direction [.advaced-
priv]#(privilege: advanced)#::
Use this parameter with the value `upgrade` to select status information
about upgrades. Use this parameter with the value `revert` to select
status information about reversions.
```

```
``[-node-status {reverting|complete|not-needed|aborted|failed|waiting|in-
progress|stopped}]`` - Node Status [.advaced-priv]#(privilege:
advanced)#::
Selects status information about upgrades or reversions that have the
status you specify on nodes where they are slated to occur. Node statuses
are:
+
  * aborted - Upgrade process aborted. Contact support personnel.
  * failed - Upgrade process failed. Contact support personnel.
  * stopped - Upgrade process stopped due to node or management
application restart. Use the xref:{relative_path}system-node-upgrade-
revert-upgrade.html[system node upgrade-revert upgrade] command to
complete the upgrade manually.
  * complete - Upgrade process completed successfully.
  * waiting - Upgrade process is waiting the replication database to come
online or for applications to be stable. If the RDB is not online, check
network connectivity using xref:{relative_path}cluster-show.html[cluster
show] and xref:{relative_path}cluster-ping-cluster.html[cluster ping-
cluster] to ensure that all nodes are healthy and in communication.



``[-node-status-msg <text>]`` - Node Status Message [.advaced-
priv]#(privilege: advanced)#::
Selects status information about upgrades or reversions that have the node
status message you specify. The node status message displays the upgrade
or reversion status of the node with which it appears. If the upgrade or
reversion fails, this message provides information that helps to diagnose
the cause of the failure.

== Examples

The following example shows typical output for a cluster with two nodes.
Status messages for each phase display information about the tasks in that
phase.
```

cluster1::*> system node upgrade-revert show

Node: node1 Status: complete

Status Message: The upgrade is complete.
Vers Phase Status Upgrade Phase Status Message
---- ---------- -------- ----------
200 pre-root applied No upgrade is required for this phase.
200 pre-apps applied Upgrade successful.
200 post-apps applied Upgrade successful.

```
Node: node2 Status: complete

Status Message: The upgrade is complete.
Vers Phase Status Upgrade Phase Status Message
---- ---------- -------- ------------------------------------------------------
200 pre-root applied No upgrade is required for this phase.
200 pre-apps applied Upgrade successful.
200 post-apps applied Upgrade successful.
6 entries were displayed.
```

```
== Related Links

* xref:{relative_path}system-node-upgrade-revert-upgrade.html[system node
upgrade-revert upgrade]
* xref:{relative_path}cluster-show.html[cluster show]
* xref:{relative_path}cluster-ping-cluster.html[cluster ping-cluster]




[[ID62ca29c5aaea120239e11de1f45346a4]]
= system node upgrade-revert upgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Run the upgrade at a specific phase.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node upgrade-revert upgrade` command manually executes an
upgrade. Use this command to execute an upgrade after issues that caused
an upgrade failure are resolved. If the upgrade is successful, no messages
display.

Before the command executes upgrades, it checks the configuration of the
nodes in the cluster. If no upgrades are needed, the command displays a
message and does not execute any upgrades.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Specifies the node that is to be upgraded. The value `local` specifies the
current node.

== Examples

This example shows command output of a node named `node0` if node
configuration is current.
```

cluster1::*> system node upgrade-revert upgrade -node node0
The node configuration is up-to-date. No upgrade is needed.

```
[[ID0e035d49de06e2d7af1d16acb310b47f]]
= system node usb-ports modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the state of the external USB ports on the next boot

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node usb-ports modify` command allows the user to change the
state of the external USB ports on the next boot.

== Parameters

``-node {<nodename>|local}`` - Node That Owns the External USB Ports
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies which node the external USB state will be
modified.

``[-is-disabled-next-boot {true|false}]`` - State of the External USB
Ports on the Next Boot [.advaced-priv]#(privilege: advanced)#::
If set to ` true ` the external USB ports on the specified node will be
disabled on the next boot.

== Examples

The following example disables the external USB ports for node1 on the
next boot:
```

cluster1::> system node usb-ports modify -node node1 -is-disabled-next-boot true
ALERT: A reboot is required for changes to take effect.

```
[[IDe740b26be6aceaa6809519fe8ff1d459]]
= system node usb-ports show
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the state of the external USB ports

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node usb-ports show` command displays the status of the
external USB ports on each node.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node That Owns the External USB Ports
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information about the
status of the external USB ports on the specified node.

``[-is-disabled {true|false}]`` - Current State of the External USB Ports
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the current state of the external USB ports.

``[-is-disabled-next-boot {true|false}]`` - State of the External USB
Ports on the Next Boot [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the state of the external USB ports after the
next boot.

``[-is-feature-supported {true|false}]`` - External USB Ports Supported?
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether or not the node even supports external
USB ports.

``[-connected-ports {true|false}]`` - Connected USB Devices [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether or not the USB ports have a device
```

```
connected to them.

== Examples

The following example displays information about the state of the external
USB ports:
```

cluster1::> system node usb-ports show
Node Disabled? Disabled on Next Boot? Supported? Devices Connected to Ports
---------- ---------- ---------- ---------- ---------- ----------
node1 false true true false

```
[[ID4c0513d9d06b089150c6cf03c5ce0632]]
= system node virtual-machine show-network-load-balancer
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display network load balancer information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node virtual-machine show-network-load-balancer` displays the
list of network load balancer probe ports for each ONTAP node in the
cluster.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - Node [.advaced-priv]#(privilege: advanced)#::
Represents the name of the ONTAP node for which information is to be
displayed. If this parameter is not specified, the command displays
information about all nodes in the cluster.
```

```
``[-vserver-name <text>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Vserver name.

``[-lif-name <text>]`` - ONTAP LIF Name [.advaced-priv]#(privilege:
advanced)#::
ONTAP logical interface name.

``[-probe-port <integer>]`` - Probe Port [.advaced-priv]#(privilege:
advanced)#::
A TCP port which is regularly probed by the network load balancer. When
the TCP port is healthy and open, the network load balancer will continue
sending traffic to an associated network route. When unhealthy, the
network load balancer will redirect all traffic intended for this route to
an alternate route.

``[-last-probe-time <MM/DD/YYYY HH:MM:SS>]`` - Last Probe Time [.advaced-
priv]#(privilege: advanced)#::
The timestamp of the most recent health probe request on this TCP port.

``[-remove-listener {true|false}]`` - Remove listener for This LIF
[.advaced-priv]#(privilege: advanced)#::
Whether or not ONTAP has programmatically told the network load balancer
to stop listening on the health probe associated with this LIF.

``[-active {true|false}]`` - Actively receiving Health Probes [.advaced-
priv]#(privilege: advanced)#::
Whether or not the network load balanacer has received a health probe
request on this TCP port, within the expected timeframe.

== Examples

The following example displays probe ports for each node in the cluster.
```

cluster1::*> system node virtual-machine show-network-load-balancer
Logical Probe
Node Vserver Interface Port Last Probe Time
--- --- --- --- ---
node1 vserver0 data_lif1 63002 5/8/2018 19:22:47
cluster1 cluster_mgmt 63001 5/8/2018 19:22:43
node2 vserver0 data_lif2 63003 5/8/2018 19:22:44
vserver0 svm_mgmt 63004 5/8/2018 19:22:50

4 entries were displayed.

```
[[ID13fbcbff1892fe41e312857931088819]]
```

```
= system node virtual-machine disk-object-store create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Define the configuration for an object store

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node virtual-machine disk-object-store create` command adds an
object store container to a node's configuration. All objects within the
container will be added as disks to the specified node.

== Parameters

``-node <nodename>`` - ONTAP Node Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the ONTAP node to which the object store container
will be added.

``-object-store-name <object store name>`` - ONTAP Name for this Object
Store Config [.advaced-priv]#(privilege: advanced)#::
Specifies the name that will be used to identify the object store
configuration.

``-server <text>`` - Fully Qualified Domain Name of the Object Store
Server [.advaced-priv]#(privilege: advanced)#::
Specifies the object store server where the container is hosted.

``-port <integer>`` - Port Number of the Object Store [.advaced-
priv]#(privilege: advanced)#::
Specifies the port number to connect to the object store server.

``-container-name <text>`` - Container Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the container to be added.

``-azure-account <text>`` - Azure Storage Account [.advaced-
priv]#(privilege: advanced)#::
Specifies the Azure storage account.

``-azure-private-key <text>`` - Azure Storage Account Access Key
```

```
[.advaced-priv]#(privilege: advanced)#::
Specifies the access key required to authenticate requests to the Azure
object store.

``[-update-partner <true>]`` - Update HA Partner [.advaced-
priv]#(privilege: advanced)#::
Specify this parameter when the system is running in an HA configuration.

== Examples

The following example adds a container to the specified node.
```

cluster1::*> system node virtual-machine disk-object-store create
-node node1 -object-store-name objstore1 -server storageaccount1.blob.core.windows.net
-container-name container1 -azure-account storageaccount1
-azure-private-key
XpSUcS/f1sl4sHfDuzYeyU3Yz9dNqVEsxDv48/P8Zk8j0uDoWYnsf/8JBhlHImH/RP9IO6maKLYqEXAMPLEKEY
== -update-partner

```
[[IDc7a2750814125a0b986589210855bab1]]
= system node virtual-machine disk-object-store delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete the configuration of an object store

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node virtual-machine disk-object-store delete` command removes
an object store container from a node's configuration.

== Parameters

``-node <nodename>`` - ONTAP Node Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the ONTAP node from which the object store container
will be removed.

``-object-store-name <object store name>`` - ONTAP Name for this Object
Store Config [.advaced-priv]#(privilege: advanced)#::
Specifies the name that will be used to identify the object store
configuration.

``[-update-partner {true|false}]`` - Update HA Partner [.advaced-
priv]#(privilege: advanced)#::
Specify this parameter when the system is running in an HA configuration.

== Examples

The following example removes a container from the specified node.
```

cluster1::*> system node virtual-machine disk-object-store delete
-node node1 -object-store-name objstore1 -update-partner

```
[[ID87bed931d45e0b674b65be0ef0ba9dc4]]
= system node virtual-machine disk-object-store modify
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the configuration of an object store

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node virtual-machine disk-object-store modify` command updates
one or more object store configuration parameters.

== Parameters

``-node <nodename>`` - ONTAP Node Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the name of the ONTAP node for which the object store
configuration will be modified.

``-object-store-name <object store name>`` - ONTAP Name for this Object
Store Config [.advaced-priv]#(privilege: advanced)#::
Specifies the name that will be used to identify the object store
configuration.

``[-server <text>]`` - Fully Qualified Domain Name of the Object Store
Server [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the new Fully Qualified Domain Name
(FQDN) of the same object store server.

``[-azure-private-key <text>]`` - Azure Storage Account Access Key
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies a new access key for the storage
account.

``[-update-partner <true>]`` - Update HA Partner [.advaced-
priv]#(privilege: advanced)#::
Specify this parameter when the system is running in an HA configuration.

== Examples

The following example updates the stored private key for an Azure
container on the specified node.
```

cluster1::*> system node virtual-machine disk-object-store modify

-node node1 -object-store-name objstore1
-azure-private-key
XpSUcS/f1sl4sHfDuzYeyU3Yz9dNqVEsxDv48/P8Zk8j0uDoWYnsf/8JBhlHImH/RP9IO6maKLYqEXAMPLEKEY
== -update-partner

```
[[ID80076649a19b60e82301ea9015e21c60]]
= system node virtual-machine disk-object-store show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of object store configurations

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system node virtual-machine disk-object-store show` command displays
the list of object store containers that contain each node's disks.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node <nodename>]`` - ONTAP Node Name [.advaced-priv]#(privilege:
advanced)#::
Represents the name of the ONTAP node for which information is to be
displayed. If this parameter is not specified, the command displays
information about all nodes in the cluster.

``[-object-store-name <object store name>]`` - ONTAP Name for this Object
Store Config [.advaced-priv]#(privilege: advanced)#::
Selects object store configurations with the specified name.

``[-server <text>]`` - Fully Qualified Domain Name of the Object Store
Server [.advaced-priv]#(privilege: advanced)#::
Selects containers on the specified server.
```

```
``[-port <integer>]`` - Port Number of the Object Store [.advaced-
priv]#(privilege: advanced)#::
Selects containers attached on the specified port.

``[-container-name <text>]`` - Container Name [.advaced-priv]#(privilege:
advanced)#::
Selects containers with the specified name.

``[-azure-account <text>]`` - Azure Storage Account [.advaced-
priv]#(privilege: advanced)#::
Selects containers in the specified Azure storage account.

``[-alive {true|false}]`` - Is Server Alive [.advaced-priv]#(privilege:
advanced)#::
Selects containers based on their aliveness state, as seen from the ONTAP
node.

== Examples

The following example displays the list of containers for each node in the
cluster.
```

cluster1::*> system node virtual-machine disk-object-store show
Object Store
Node Name Azure Storage Account Container Name Alive
---------- ---------- ---------- ---------- -----
node1
objstore1 storageaccount1 container1 true
node2
objstore1 storageaccount1 container1 true
2 entries were displayed.

```
[[IDbe675612c39495bd0e1a522ff2db0adb]]
= system node virtual-machine hypervisor show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display hypervisor information about Data ONTAP-v nodes

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The `system node virtual-machine hypervisor show` command displays
information for each hypervisor that is hosting a ONTAP virtual machine.
The output contains the hypervisor-specific information, such as host name
and IP address, as well as network configuration details. The command only
scans hypervisors on which ONTAP virtual machines are installed. To filter
command output, specify any number of optional fields listed below.

== 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>|local}]`` - Node::
The name of the ONTAP node running in a virtual machine for which
information is to be displayed. If this optional parameter is not
specified, the command displays information about all nodes in the
cluster.

``[-vm-uuid <UUID>]`` - UUID of the Virtual Machine::
The hypervisor-supplied unique ID for this virtual machine. This optional
parameter selects information about the hypervisor on which the ONTAP
virtual machine is running with the specified UUID. Since UUID is unique
per host, an alternative and easier way is to use `-node` to filter out
the same information.

``[-vmhost-bios-release-date <text>]`` - Release Date for the Hypervisor
BIOS::
The release date for the currently running hypervisor BIOS. This optional
parameter selects information about the hypervisors that have the
specified BIOS release date.

``[-vmhost-bios-version <text>]`` - Current BIOS Version of the Hypervisor
Physical Chassis::
The current BIOS version of the hypervisor physical chassis. This optional
parameter selects information about the hypervisors that are running with
the specified BIOS version.

``[-vmhost-boot-time <text>]`` - Time When Hypervisor was Last Booted::
The time when the hypervisor was last booted. This optional parameter
```

```
selects information about the hypervisors which were last booted at the
specified boot time.

``[-vmhost-cpu-clock-rate <integer>]`` - Speed of the Hypervisor CPU Cores
(MHz)::
The speed of the hypervisor CPU cores. This optional parameter selects
information about the hypervisors that are running with the specified CPU
clock rate.

``[-vmhost-cpu-core-count <integer>]`` - Number of Physical CPU Cores on
the Hypervisor::
The number of physical CPU cores on the hypervisor. Physical CPU cores are
the processors contained by a CPU package. This optional parameter selects
information about the hypervisors that are running with the specified CPU
cores.

``[-vmhost-cpu-socket-count <integer>]`` - Number of Physical CPU Packages
on the Hypervisor::
The number of physical CPU packages on the hypervisor. Physical CPU
packages are chips that contain one or more processors. Processors
contained by a package are also known as CPU cores. For example, one dual-
core package consists of one chip that contains two CPU cores. This
optional parameter selects information about the hypervisors that are
running with the specified CPU sockets.

``[-vmhost-cpu-thread-count <integer>]`` - Number of Physical CPU Threads
on the Hypervisor::
The number of physical CPU threads on the hypervisor. This optional
parameter selects information about the hypervisors that are running with
the specified CPU threads.

``[-vmhost-gateway <text>]`` - Default Gateway for the Hypervisor::
The default gateway for the hypervisor. This optional parameter selects
information about the hypervisors with the specified gateway address.

``[-vmhost-hardware-vendor <text>]`` - Hardware Vendor of the Hypervisor::
The name of hypervisor hardware manufacturer. This optional parameter
selects information about the hypervisors with the specified hardware
vendor.

``[-vmhost-hypervisor <text>]`` - Complete Product Name, including the
Version Information for the Hypervisor::
The complete product name, including the version information for the
hypervisor. This optional parameter selects information about the
hypervisors that are running with the specified hypervisor version.
```

```
``[-vmhost-ip-address <text>]`` - Primary IP Address Assigned to the
Hypervisor::
The primary IP address assigned to the hypervisor. This optional parameter
selects information about the hypervisors with the specified IP address.

``[-vmhost-memory <integer>]`` - Physical Memory Size of the Hypervisor
(Bytes)::
The physical memory size of the hypervisor in bytes. This optional
parameter selects information about the hypervisors that are running with
the specified physical memory.

``[-vmhost-model <text>]`` - Hypervisor Manufacturer-Supplied Hardware
Model Name::
The hypervisor manufacturer-supplied hardware model name. This optional
parameter selects information about the hypervisors with the specified
hardware model.

``[-vmhost-name <text>]`` - Hostname of the Hypervisor::
The host name assigned to the hypervisor. This optional parameter selects
information about the hypervisor with the specified host name.

``[-vmhost-netmask <text>]`` - Subnet Mask Address for the Hypervisor::
The subnet mask address for the hypervisor. This optional parameter
selects information about the hypervisors with the specified netmask
address.

``[-vmhost-processor-id <text>]`` - Processor ID of the Hypervisor::
The processor ID of the hypervisor. This optional parameter selects
information about the hypervisors with the specified processor ID.

``[-vmhost-processor-type <text>]`` - CPU Model of the Hypervisor::
The CPU model of the hypervisor. This optional parameter selects
information about the hypervisors that are running with the specified
processor type.

``[-vmhost-software-vendor <text>]`` - Name of the Virtual Machine
Software Manufacturer::
The name of the virtual machine software manufacturer. This optional
parameter selects information about the hypervisors with the specified
software vendor.

``[-vmhost-uuid <UUID>]`` - UUID of the Hypervisor::
A unique ID for the hypervisor. This optional parameter selects
information about the hypervisor with the specified UUID.

``[-vmhost-error <text>]`` - Error in case Hypervisor Info Retrieval
Fails::
```

```
Displays a list of nodes on which the hypervisor has received the
specified error. This parameter is most useful when entered with
wildcards.

``[-vm-custom-max-capacity <integer>]`` - Maximum Storage Capacity of the
Virtual Machine (in TB)::
The maximum system capacity (in TB) that can be configured on the VM. This
optional parameter selects information about the node's storage capacity.

== Examples

The following example shows typical output from the `system node virtual-
machine hypervisor show` command for the ONTAP virtual machines running in
the cluster.
```

cluster1::> system node virtual-machine hypervisor show

Virtual Machine Info

Node: node1
VM UUID: 123abcde-4f5g-6h78-i9j0-k12l3m4567np
Hypervisor Info

Hardware Vendor: VMware, Inc.
Model: VMware Virtual Platform
Software Vendor: Unknown
Hypervisor: VMware ESX 4.1.0 build-12345
Host Name: myesx.example.com
Last Boot Time: 2014-01-01T01:23:45.678901-23:45
Host UUID: 00000000-0000-0000-0000-0012a3456789
BIOS Version: S1234.5.6.7.8.901234567890
BIOS Release Date: 2013-01-01T00:00:00Z
CPU Packages: 2
CPU Cores: 12
CPU Threads: 24
Processor ID: 0000:0000:0000:0010:0010:0100:1100:0010
Processor Type: Intel® Xeon® CPU X5670 @ 2.93GHz
CPU MHz: 2925
Memory Size: 4227858432
IPv4 Configuration: IP Address: 192.168.0.1
Netmask: 255.255.255.0
Gateway: 192.165.0.1
Virtual Machine Info

Node: node2
VM UUID: 123abcde-4f5g-6h78-i9j0-k98l7m6543yz
Hypervisor Info

Hardware Vendor: VMware, Inc.
Model: VMware Virtual Platform
Software Vendor: Unknown

Error: ServerFaultCode:
InvalidLoginFault type='InvalidLogin'

2 entries were displayed.

```
[[ID1c7805fedd059c1bc65ff1f94fcfb29a]]
= system node virtual-machine instance show-system-disks
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about virtual machine system disks

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system node virtual-machine instance show-system-disks` command
displays information about the system disks (non-data disks) attached to
the virtual machine. Data disk information is available using the command
xref:{relative_path}storage-disk-show.html[storage disk show]`-virtual-
machine-disk-info` .

== 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>|local}]`` - Node Name::
Selects disk information for nodes that match this parameter.

``[-vmdisk-purpose <text>]`` - Purpose of the System Disk::
Selects disk information for disks that match this parameter. Values
include: ``_Boot_`` , ``_NVRAM_`` , and ``_Core_`` .

``[-vmdisk-type <text>]`` - Type of the System Disk::
Selects disk information for disks that match this parameter. Possible
values are: ``_VMDISK_`` , ``_SSD_`` .
```

```
``[-vmdisk-name <text>]`` - System Disk Name::
Selects disk information for disks that match this parameter. The virtual
machine ID of the system disk.

``[-vmdisk-capacity <text>]`` - Size of the System Disk (GB)::
Selects disk information for disks that match this parameter. The size of
the system disk.

``[-vmdisk-file-name <text>]`` - File Name of the System Disk Used By the
Hypervisor::
Selects disk information for disks that match this parameter. The virtual
machine file name of the disk. Each system disk is mapped to a unique VM
disk file.

== Examples

The following example shows typical output from the command.
```

cluster1::> system node virtual-machine instance show-system-disks
Disk Disk Disk Disk
Node Purpose Type Name Capacity GB VM Disk File Name
---- ------- ------- ---- ---------- ----------
node1
boot SSD da0 10 node1-vm-disk-boot
nvram SSD da1 500 node1-vm-disk-nvram
core VMDISK da2 216 node1-vm-disk-core

```
== Related Links

* xref:{relative_path}storage-disk-show.html[storage disk show]




[[ID52ceb673d16933e9ab2e106dbc39d3bd]]
= system node virtual-machine instance show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display virtual machine instance information per node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `system node virtual-machine instance show` command displays virtual
machine information. With this information you can determine the
relationship between a ONTAP ``_node_`` and its associated virtual machine
instance running within a cloud provider. Several other details about the
virtual machine can be extracted as well, such as the cloud provider
account ID to which it belongs. To filter command output, specify any
number of optional fields listed below.

== 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>|local}]`` - Node::
This optional parameter represents the name of the ONTAP node running in a
virtual machine for which information is to be displayed. If this
parameter is not specified, the command displays information about all
nodes in the cluster.

``[-instance-id <text>]`` - ID of This Instance::
Selects the nodes that match this parameter value. A cloud provider-
supplied unique instance ID for this virtual machine, for example "i-
a9d42f89" or "db00a7755a5e4e8a8fe4b19bc3b330c3.1".

``[-account-id <text>]`` - ID of This Account::
Selects the nodes that match this parameter value. The cloud provider-
associated account ID for this virtual machine. This parameter is usually
associated with a cloud provider login ID and password.

``[-image-id <text>]`` - ID Of the Image in Use on This Instance::
Selects the nodes that match this parameter value. The image ID installed
on this virtual machine instance. It identifies a pre-defined template of
a computing device's software environment. It contains the operating
system and can also include application software, such as database
servers, middleware, and web servers. In this case, the ID refers to an
image that contains everything required to run ONTAP in the cloud.
```

```
``[-instance-type <text>]`` - Specifies System Attributes and Use Cost::
Selects the nodes that match this parameter value. A specification (as
defined by the cloud provider) that defines the memory, CPU, storage
capacity and usage cost for a virtual machine instance. Some instance
types are designed for standard applications, whereas others are designed
for CPU-intensive or memory-intensive applications and so on.

``[-region <text>]`` - Set of Resources in the Same Geographic Area::
Selects the nodes that match this parameter value. A named set of
resources in the same geographical area. For example "us-east-1" might be
the name for a collection of compute and storage resources on the eastern
coast of the United States. Typically, a region contains multiple
availability zones.

``[-version <text>]`` - Instance Metadata Version of This VM::
Selects the nodes that match this parameter value. The instance metadata
version of this virtual machine instance.

``[-availability-zone <text>]`` - Distinct Location within a Region::
Selects the nodes that match this parameter value. A distinct location
within a ``_region_`` that is insulated from failures in other
availability zones. It provides low-latency network connectivity to other
availability zones in the same region.

``[-primary-ip <text>]`` - Primary IP Address Assigned to this Instance::
Selects the nodes that match this parameter value. The primary IP address
assigned to this virtual machine instance.

``[-fault-domain <integer>]`` - Fault Domain of This Instance::
Selects the nodes that match this parameter value. A cloud provider-
assigned numerical fault domain ID for this virtual machine.

``[-update-domain <integer>]`` - Update Domain of This Instance::
Selects the nodes that match this parameter value. A cloud provider-
assigned numerical update domain ID for this virtual machine.

``[-provider <text>]`` - Provider on which this instance is running.::
Selects the nodes that match this parameter value. The provider on which
this instance is running.

``[-offer <text>]`` - Marketplace Offer for This Image::
Selects the nodes that match this parameter value. A Marketplace is an
online store that offers applications and services either built on or
designed to integrate with a particular cloud provider. A virtual machine
offer corresponds to a class of product offering from a publisher. An
offer is a collection of SKUs. Every offer appears as its own entity in
```

```
the provider Marketplace.

``[-sku <text>]`` - Marketplace SKU for This Image::
Selects the nodes that match this parameter value. A SKU is the smallest
purchasable unit of an offer. While within the same product class (offer),
SKUs allow you to differentiate between different features supported, VM
image types and billing models supported.

``[-sku-version <text>]`` - Marketplace Version of a SKU::
Selects the nodes that match this parameter value. The version for this
virtual machine SKU.

``[-resource-group-name <text>]`` - Resource Group Name of This Instance::
Selects the nodes that match this parameter value. Resource group for the
instance.

``[-cpu-platform <text>]`` - CPU Platform of the Instance::
Selects the nodes that match this parameter value. GCP only. CPU platform
of the hypervisor host. Example: ``_Intel Broadwell_`` .

``[-consumer <text>]`` - Consumer of this Instance::
Selects the nodes that match this parameter value. Consumer is based on
the platform/infrastructure on which ONTAP is launched. Example:
``_Marketplace_`` .

``[-total-network-bandwidth-limit <integer>]`` - Total Network Bandwidth
Limit in MB/s::
Selects the nodes that match this parameter value. Represents the total
network bandwidth limit for the given VM in MBps. For IBM Cloud, the limit
is calculated separately for ingress and egress. Example: ``_1024_`` .

``[-total-volume-bandwidth-limit <integer>]`` - Total Volume Bandwidth
Limit in MB/s::
Selects the nodes that match this parameter value. Represents the total
volume bandwidth limit for the given VM in MBps. For IBM Cloud, the limit
is calculated separately for ingress and egress. Example: ``_1024_`` .

``[-az-id <text>]`` - Id of Availability Zone in which the instance is
running::
Id of Availability Zone in which the instance is running

``[-storage-ontap-instance-id <text>]`` - Storage ontap instance ID for
the deployment::
Storage ONTAP instance ID for the deployment

``[-vpc-partner-api-endpoint <text>]`` - VPC partner api endpoint for the
deployment::
```

```
VPC partner api endpoint for the deployment

``[-self-setup {true|false}]`` - Deployed and configured without BlueXP::
Deployed and configured without BlueXP

== Examples

The following examples illustrate typical output from the `system node
virtual-machine instance show` command for a virtual machine running in a
cloud provider environment.
```

cluster1::> system node virtual-machine instance show
Node: node1
Instance ID: i-b9c42e97
Account ID: 751083215869
Image ID: ami-7fb4a1c6
Instance Type: m3.xlarge
Region: us-east-1
Metadata Version: 2010-08-31
Availability Zone: us-east-1d
Primary IP: 192.168.0.1
Provider: AWS
Consumer: Marketplace
Storage Type: SSD
IOPS: 3000
Storage Capacity (GB): 1024
Throughput Capacity (MBps): 1000

cluster1::> system node virtual-machine instance show
Node: node1
Instance ID: 090556da-d4fa-764f-a9f1-63614eda019a
Metadata Version: 2012-11-30
Fault Domain: 0
Update Doamin: 0
Primary IP: 192.168.0.1
Provider: Azure
Offer: netapp-ontap-cloud
SKU: ontap_cloud_pgo_sn
SKU Version: 9.4.20180510
Resource Group Name: resourcegroup1
Account ID: 228e471c-3b42-4ae7-9b59-df5bb5e6228d

cluster1::> system node virtual-machine instance show
Node: node1
Instance ID: 1234567890123456789
Account ID: customer-project-id
Image ID: projects/project-id/global/images/image-id
Instance Type: n1-standard-8
CPU Platform: Intel Broadwell
Region: us-east4
Metadata Version: v1
Availability Zone: us-east4-b

Primary IP: 192.168.0.1
Provider: GCP

cluster1::> system node virtual-machine instance show
Node: node1
Instance ID: 0757_c00488fc-1de9-45e2-b895-4b05119cece2
Image ID: r014-31d63c06-b5da-4147-8751-c2502dfad35a
Instance Type: bx2-8x32
CPU Platform: Intel Xeon Processor (Cascadelake)
Metadata Version: 2022-05-31
Availability Zone: us-east-1
Primary IP: 10.160.96.82
Provider: IBM Cloud
Consumer: -
Resource Group Name: unknown
Total Network Bandwidth limit (MBps): 500
Total Volume Bandwidth limit (MBps): 1500
Storage Ontap Instance ID: b27115df-1f54-4ee6-85a6-1ae04a7fe42f

```
:leveloffset: -1


= system script commands

:leveloffset: +1


[[IDbaa79ff9e43aa618543b148171c7e326]]
= system script delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete saved CLI session logs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system script delete` command deletes files that contain CLI session
records. Use the xref:{relative_path}system-script-show.html[system script
show] command to display saved CLI sessions.

== Parameters

``-username <text>`` - Log Owner Username::
Use this parameter to specify the name of the user whose CLI session
record files are deleted. The default is the username is that of the
logged in user.

``-filename <text>`` - Log Filename::
Use this parameter to specify the names of CLI session record files to
delete.

== Examples

The following example shows how to delete the files named `sessionlog2`
and `sessionlog3` .
```

cluster1::> system script delete -filename sessionlog2,sessionlog3

```
The following example deletes all saved script files.
```

cluster1::> system script delete *

```
== Related Links

* xref:{relative_path}system-script-show.html[system script show]




[[ID69dcddfa154311f8aa0f04b8100477d0]]
= system script show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display saved CLI session logs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system script show` command displays information about files that
contain records of CLI sessions.

For security reasons, the command normally displays only the script files
created by the logged in user. Administrative users can display all log
files using the ``_-user_`` parameter.

== 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.

| ``[-user  ]``::
Use this parameter to display all script files created by all users, along
with the username associated with each file.

| ``[-instance  ]`` }::
```

```
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-username <text>]`` - Log Owner Username::
Use this parameter to display information only about files saved by the
user you specify. The default username is that of the logged in user.

``[-filename <text>]`` - Log Filename::
Use this parameter to display information only about files that have the
file name you specify.

``[-size-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Logfile Size Limit::
Use this parameter to display information only about files that have the
size limit you specify.

``[-state <State of CLI session log>]`` - Current State::
Use this parameter to display information only about files that have the
state you specify. Valid values for this parameter are `open-and-logging`
, `file-full` , and `file-closed` .

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Current Logfile Size::
Use this parameter to display information only about files that are the
size you specify.

``[-mtime <MM/DD/YYYY HH:MM:SS>]`` - Last Modification Time::
Use this parameter to display information only about files that were last
modified at the date and time you specify.

``[-this-session {yes|no}]`` - Session is Logging::
Use this parameter with the value `yes` to display information only about
files that are recording the current CLI session. Use this parameter with
the value `no` to display information only about files that are not
recording the current CLI session.

== Examples

The following example displays typical system script information.
```

cluster1::> system script show
This
FileName Sess State Size Last Mod Date
--- --- --- --- ---
sessionlog1 no file-closed 435B 12/2/2008 10:51:12
sessionlog2 yes open-and-logging 193B 12/2/2008 10:51:29
2 entries were displayed.

```
[[ID89df469e5891c2e1f98df8e160f50672]]
= system script start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start logging all CLI I/O to session log

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system script start` command starts creating a record of your CLI
session. The record is stored in a file. Use the
xref:{relative_path}system-script-show.html[system script show] ``_-this-
session_```yes` command to display files that are recording the current
CLI session. Use the xref:{relative_path}system-script-stop.html[system
script stop] command to stop recording the current CLI session.

== Parameters

``-filename <text>`` - Filename to Log To::
Use this parameter to specify the file name to which the CLI session
record is saved.

``-size-limit {<integer>[KB|MB|GB|TB|PB]}`` - Logfile Size Limit Max:2GB::
Use this parameter to specify the maximum size of the file that contains
the CLI session record. When the file size reaches this limit, recording
stops. The default file size limit is `1 MB` . The maximum file size limit
is `2 GB` .

== Examples

The following example shows how to start creating a record of the CLI
session in a file named `sessionlog3` . The size limit of this file is `20
MB` .
```

cluster1::> system script start -filename sessionlog3 -size-limit 20MB

```
== Related Links

* xref:{relative_path}system-script-show.html[system script show]
* xref:{relative_path}system-script-stop.html[system script stop]




[[IDfa66652dc54789397be7863dc54394eb]]
= system script stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stops logging CLI I/O

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system script stop` command stops creating a record of your CLI
session, if you started creating the record by using the
xref:{relative_path}system-script-start.html[system script start] command.
Use the xref:{relative_path}system-script-show.html[system script show]
``_-this-session_```yes` command to display files that are recording the
current CLI session.

== Examples

The following example shows how to stop creating a record of your CLI
session.
```

cluster1::> system script stop

```
== Related Links

* xref:{relative_path}system-script-start.html[system script start]
* xref:{relative_path}system-script-show.html[system script show]
```

```
[[ID0730e5255357fb9d535544dc2357d3c5]]
= system script upload
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Upload the selected CLI session log

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system script upload` command uploads a CLI session record file to a
remote location. Specify the remote location using an FTP or HTTP URI. Use
the xref:{relative_path}system-script-show.html[system script show]
command to display saved CLI sessions. Use the xref:{relative_path}system-
script-start.html[system script start] command to record a CLI session and
save it to a file.

== Parameters

``-username <text>`` - Username If Not Your Own::
Use this parameter to specify the name of the user who owns the file to
upload. By default, this is the user who is logged in.

``-filename <text>`` - Filename to Log To::
Use this parameter to specify the name of a file to be uploaded.

``-destination {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}``
- URI to Send File To::
Use this parameter to specify the FTP or HTTP destination of the file.

== Examples

The following example shows how to upload the file named `sessionlog3` to
the destination `ftp://now.example.com/cli_sessions` .
```

cluster1::> system script upload -filename sessionlog3 -destination ftp://now.example.com/cli_sessions

```
== Related Links

* xref:{relative_path}system-script-show.html[system script show]
* xref:{relative_path}system-script-start.html[system script start]
```

```
:leveloffset: -1


= system service-processor commands

:leveloffset: +1


[[ID9dcec64e4a269cd72fe0aa11b372ef05]]
= system service-processor reboot-sp
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reboot the Service Processor on a node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor reboot-sp` command reboots the Service
Processor of the specified node.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node whose Service Processor is to be
rebooted.

``[-image {primary|backup}]`` - Image to Boot with After Reboot::
This parameter specifies the image that the Service Processor uses after
the reboot. By default, the `primary` image is used. Avoid booting the SP
from the backup image. Booting from the backup image is reserved for
troubleshooting and recovery purposes only. It might require that the SP
automatic firmware update be disabled, which is not a recommended setting.
You should contact technical support before attempting to boot the SP from
the backup image.

== Examples
```

```
The following command reboots the Service Processor of node "node1" into
the `primary` image.
```

cluster1::> system service-processor reboot-sp -node node1 -image primary

NOTE : If your console connection is through the SP, it will be disconnected.
Do you want to reboot the SP ? {y|n}: y

cluster1::>

```
The following command reboots the Service Processors of all nodes. Since
`-image` is not specified, the Service Processors will boot into the
`primary` image.
```

cluster1::> system service-processor reboot-sp -node *

NOTE : If your console connection is through the SP, it will be disconnected.
Do you want to reboot the SP ? {y|n}: y
2 entries were acted on.

cluster1::>

```
[[IDdbfdb542180cb054fee292212b3f515f]]
= system service-processor show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Service Processor information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor show` command displays information about the
Service Processor of each node in a cluster. You can limit output to
specific types of information and specific nodes in the cluster, or filter
output by specific field values.

In case a node is offline or its Service Processor management daemon is
down, the command displays the last known IP address of its Service
Processor. Only the IP address is displayed in such cases.
```

```
== 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>|local}]`` - Node::
Selects information for the Service Processor of the specified node.

``[-type {SP|NONE|BMC}]`` - Type of Device::
Selects information for the Service Processors of the specified type.

``[-status {online|offline|sp-daemon-offline|node-
offline|degraded|rebooting|unknown|updating}]`` - Status::
Selects information for the Service Processors whose status matches the
specified value.

``[-ip-configured {true|false}]`` - Is Network Configured::
Selects information for the Service Processors whose network is configured
(`true` ) or not configured (`false` ).

``[-address <IP Address>,...]`` - Public IP Address::
Selects information for the Service Processors that use the specified IP
address or addresses.

``[-mac <MAC Address>]`` - MAC Address::
Selects information for the Service Processors that use the specified MAC
address.

``[-fw-version <text>]`` - Firmware Version::
Selects information for the Service Processors that are running the
specified firmware version.

``[-part-num <text>]`` - Part Number::
Selects information for the Service Processors that have the specified
part number.

``[-serial-num <text>]`` - Serial Number::
Selects information for the Service Processors that have the specified
serial number.

``[-dev-rev <text>]`` - Device Revision::
```

```
Selects information for the Service Processors that have the specified
device revision.

``[-autoupdate-enabled {true|false}]`` - Is Firmware Autoupdate Enabled::
Selects information for the Service Processors that have the specified
status for firmware automatic update.

== Examples

The following example displays basic information for the Service
Processors of all the nodes.
```

cluster1::> system service-processor show
IP Firmware
Node Type Status Configured Version IP Address
---------- ---- ---------- ---------- -------- ----------
node1 SP online true 2.2X5 192.168.1.201
node2 SP online true 2.2X5 192.168.1.202
2 entries were displayed.

cluster1::>

```
The following example displays all available information for the Service
Processors of all the nodes.
```

cluster1::> system service-processor show -instance
Node: node1
Type of Device: SP
Status: online
Is Network Configured: true
Public IP Address: 192.168.1.201
MAC Address: ab:cd:ef:fe:ed:01
Firmware Version: 2.2X5
Part Number: Not Applicable
Serial Number: Not Applicable
Device Revision: Not Applicable
Is Firmware Autoupdate Enabled: true
Node: node2
Type of Device: SP
Status: online
Is Network Configured: true
Public IP Address: 192.168.1.202
MAC Address: ab:cd:ef:fe:ed:02
Firmware Version: 2.2X5
Part Number: Not Applicable
Serial Number: Not Applicable
Device Revision: Not Applicable
Is Firmware Autoupdate Enabled: true
2 entries were displayed.

cluster1::>

```
The following example displays only the type, status and firmware version
for the Service Processors of all the nodes.
```

cluster1::> system service-processor show -fields type,status,fw-version
node type status fw-version
---------- ---- ------ ----------
node1 SP online 2.2X5
node2 SP online 2.2X5
2 entries were displayed.

cluster1::>

```
[[ID0f662875bfb6b40bd1fa38c9ba5a2fee]]
= system service-processor api-service check
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Check API Service availability in SP/BMC

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system service-processor api-service check` checks the availability
of the SP API and web service on Service Processors (SP) or Baseboard
Management Controllers (BMC) in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node to initiate the check operation. By
default, all nodes check availability.

``[-port <integer>]`` - SP/BMC API Service Port [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the port number on the SP or BMC used for the API
service. By default, port 50000 is used.

``[-target-node {<nodename>|local}]`` - Target Node [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the target node whose SP or BMC API and web
service is to be checked. By default, all node SPs or BMCs are checked.

== Examples

The following example checks the availability of the API and web service
on all node SPs or BMCs.
```

cluster1::*> sp api-service check -node *

Date: Tue Oct 27 03:09:19 2020
Node: cluster1-01
Port: 50000

```
Local Private IPv4: ok (192.0.2.82)
Local Public IPv4: ok (10.233.4.22)
Local Public IPv6: ok (2020::2023)
Target Node: cluster1-02
IPv4: ok (10.233.4.24)
IPv6: ok (2020::2022)

Port: 443
Local Private IPv4: ok (192.0.2.82)
Local Public IPv4: ok (10.233.4.22)
Local Public IPv6: ok (2020::2023)
Target Node: cluster1-02
IPv4: ok (10.233.4.24)
IPv6: ok (2020::2022)

Date: Tue Oct 27 03:09:19 2020
Node: cluster1-02
Port: 50000
Local Private IPv4: ok (192.0.2.81)
Local Public IPv4: ok (10.233.4.24)
Local Public IPv6: ok (2020::2022)
Target Node: cluster1-01
IPv4: ok (10.233.4.22)
IPv6: ok (2020::2023)

Port: 443
Local Private IPv4: ok (192.0.2.81)
Local Public IPv4: ok (10.233.4.24)
Local Public IPv6: ok (2020::2022)
Target Node: cluster1-01
IPv4: ok (10.233.4.22)
IPv6: ok (2020::2023)
2 entries were acted on.

cluster1::*>
```

```
[[IDe7989ab656933db6c49c1beb6a9d433a]]
= system service-processor api-service disable-installed-certificates
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable user-installed certificates for the service processor API service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command disables user-installed certificates for secure communication
with the service processor API service. Default certificates are then
auto-generated.

== Examples

The following example disables user-installed certificates for the service
processor API service.
```

cluster1::> system service-processor api-service disable-installed-certificates

```
[[ID313862ccb8efab2482ff24555a532b18]]
= system service-processor api-service enable-installed-certificates
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable user-installed certificates for the service processor API service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command enables user-installed certificates for secure communication
with the service processor. Use the xref:{relative_path}security-
certificate-install.html[security certificate install] command to install
client, server and CA certificates.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Use this parameter to specify the Vserver on which certificates are
installed.

``-server-cert <text>`` - Name of the Server Certificate::
Use this parameter to specify the unique name of the server certificate.

``-client-cert <text>`` - Name of the Client Certificate::
Use this parameter to specify the unique name of the client certificate.

``-rootca-cert <text>`` - Names of the Root CA Certificates::
Use this parameter to specify the unique names of server-ca or client-ca
certificate.

== Examples

The following example installs server, client and rootca certificates and
then enables those certificates for secure communication with the service
processor.
```

cluster1::> security certificate install -vserver cluster1 -type server
cluster1::> security certificate install -vserver cluster1 -type client
cluster1::> security certificate install -vserver cluster1 -type server-ca
cluster1::> security certificate show-user-installed

Vserver Serial Number Certificate Name Type
--- --- --- ---
cluster1 1533F133482E800F
xxx-ca server-ca
Certificate Authority: xxx-ca
Expiration Date: Sat Jun 01 05:11:41 2019

cluster1 1533F273AA311FDB
xxx-client client
Certificate Authority: xxx-ca
Expiration Date: Fri May 31 05:34:37 2019

cluster1 1533F1B321E55242
xxx-server server
Certificate Authority: xxx-ca
Expiration Date: Fri May 31 05:20:50 2019

cluster1::> system service-processor api-service enable-installed-certificates -vserver cluster1 -server-cert xxx-server -client-cert xxx-client -rootca-cert xxx-ca

```
== Related Links

* xref:{relative_path}security-certificate-install.html[security
certificate install]




[[ID956d68d220f1dd946179c8e1f9bccea7]]
= system service-processor api-service modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify service processor API service configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system service-processor api-service modify` command modifies the
Service Processor (SP) or Baseboard Management Controller (BMC) API
service and web service configuration. The SP API and web is a secure
network API that enables ONTAP to communicate with the SP or BMC over the
network.
```

```
== Parameters

{ ``[-sp-api-enabled {true|false}]`` - Is SP API Service Enabled
[.advaced-priv]#(privilege: advanced)#::
This parameter enables or disables the SP API service of the SP or BMC.
When the SP API service is disabled, features like network-based firmware
updates and network-based down node log collection will not be available,
and the slower serial-interface will be used for firmware updates and down
node log collection.

``[-sp-api-port <integer>]`` - SP API Service Port [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the port number on the SP or BMC used for the SP
API service. By default, port 50000 is used.

``[-sp-api-limit-access {true|false}]`` - Only Cluster Members Can Access
the API Service [.advaced-priv]#(privilege: advanced)#::
This parameter restricts SP API service access to cluster members only. By
default, it is enabled.

| ``[-web-enabled {true|false}]`` - Is SP Web Service Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter enables or disables the web service of the SP or BMC. When
the web service is disabled, features like network-based firmware updates
and network-based down node log collection will not be available, and the
slower serial-interface will be used for firmware updates and down node
log collection.

``[-web-limit-access {true|false}]`` - Only Cluster Members Can Access the
Web Service [.advaced-priv]#(privilege: advanced)# }::
This parameter restricts web service access to cluster members only. By
default, it is enabled.

== Examples

The following example modifies the port number used for the SP API service
and then disables the SP API service.
```

cluster1::*>system service-processor api-service modify -sp-api-enabled true -sp-api-port 50001 -sp-api-limit -access true

cluster1::*>system service-processor api-service show
Service Processor API service configuration
SP API Enabled: true
SP API Port: 50001
SP API Limit Access: true
Web Enabled: true

Web Limit Access: true
Server Certificate: -internal-
Client Certificate: -internal-
CA Certificate: -internal-

```
```

cluster1::*>system service-processor api-service modify -sp-api-enabled false -sp-api-port 50001 -sp-api-limit -access true

cluster1::*>system service-processor api-service show
Service Processor API service configuration
SP API Enabled: false
SP API Port: 50001
SP API Limit Access: true
Web Enabled: true
Web Limit Access: true
Server Certificate: -internal-
Client Certificate: -internal-
CA Certificate: -internal-

```
The following example disables the SP Web service.
```

cluster1::*>system service-processor api-service modify -web-enabled false -web-limit-access true

cluster1::*>system service-processor api-service show
Service Processor API service configuration
SP API Enabled: false
SP API Port: 50001
SP API Limit Access: true
Web Enabled: false
Web Limit Access: true
Server Certificate: -internal-
Client Certificate: -internal-
CA Certificate: -internal-

```
[[IDd93f832105b7006f4ef5d51e8346ef86]]
= system service-processor api-service regenerate-ssh-auth-key
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Regenerate SSH Auth Private and Public Key

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor api-service regenerate-ssh-auth-service`
command regenerates public and private keys for SSH public key
authentication between ONTAP nodes and their service processor (SP) or
baseboard management controller (BMC).

== Parameters

``-key-type {ecdsa-256|ecdsa-384|rsa-2048|rsa-3076|rsa-4096}`` - SSH Key
Type::
This parameter specifies the SSH authentication key type and key size that
must be used when regenerating the keys. Examples: ecdsa-256, rsa-2048.

== Examples

The following example regenerates keys based on key type ECDSA and key
size 256.
```

cluster1::> system service-processor api-service regenerate-ssh-auth-key -key_type ecdsa-256

```
[[ID7790052a7dec7a361abfbe152fdaf2f2]]
= system service-processor api-service renew-internal-certificates
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Renew SSL and SSH certificates used for secure communication with the
service processor

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `system service-processor api-service renew-internal-certificates `
command generates the certificates used for secure communication with the
service processor API service. This command is not allowed if user-
installed certificates are enabled.

== Examples

The following example generates new default host and root-ca certificates.
```

cluster1::*> system service-processor api-service renew-internal-certificates

```
[[ID81af2d8fd4f5b83a7676f53db3fe75a5]]
= system service-processor api-service show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display service processor API service configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system service-processor api-service show` command displays the
Service Processor (SP) or Baseboard Management Controller (BMC) API and
web service configuration.

== Examples

The following example displays the SP or BMC API and web service
configuration:
```

cluster1::*> system service-processor api-service show
Service Processor API service configuration
SP API Enabled: true
SP API Port: 50000
SP API Limit Access: true
Web Enabled: true
Web Limit Access: true
Server Certificate: -internal-
Client Certificate: -internal-
CA Certificate: -internal-

```
[[ID088f853b365fc53b8b87b97b7a55bbc3]]
= system service-processor image modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable/Disable automatic firmware update

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor image modify` command enables or disables
automatic firmware update on the Service Processor of specified node or
nodes.

== Parameters

``-node {<nodename>|local}`` - Node::
The parameter specifies the node on which automatic Service Processor
firmware update is to be enabled or disabled.

``[-autoupdate {true|false}]`` - Firmware Autoupdate::
Setting this parameter to `true` enables automatic firmware update.
Setting this parameter to `false` disables automatic firmware update. This
is a mandatory parameter.

== Examples

The following command enables automatic firmware update for the Service
Processor on the local node.
```

cluster1::> system service-processor image modify -node local -autoupdate true

```
The following command enables automatic firmware update for the Service
Processors on all the nodes.
```

cluster1::> system service-processor image modify -node * -autoupdate true
2 entries were modified.

```
[[IDf1b22bac019dafa83d4b8d8ba175379d]]
```

```
= system service-processor image show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the details of currently installed Service Processor firmware
image

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor image show` command displays information
about the currently installed firmware images on the Service Processor of
each node in a cluster. You can limit output to specific types of
information and specific nodes in the cluster, or filter output by
specific field values.

== 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>|local}]`` - Node::
Selects firmware image information for the Service Processor of the
specified node.

``[-image {primary|backup}]`` - Image::
Selects firmware image information for the Service Processors that are
running the primary or backup image as specified.

``[-type {SP|NONE|BMC}]`` - Type::
 Selects firmware image information for the Service Processors of the
specified type.

``[-status {installed|corrupt|updating|auto-updating|none}]`` - Image
Status::
Selects firmware image information for the Service Processors whose image
```

```
status matches the specified value.

``[-is-current {true|false}]`` - Is Image Current::
Selects firmware image information for the SP whose current image matches
the specified status. This parameter indicates the partition (primary or
backup) that the SP is currently booted from, not whether the installed
firmware version is most current.

``[-version <text>]`` - Firmware Version::
Selects firmware image information for the Service Processors running the
specified firmware version.

``[-autoupdate {true|false}]`` - Firmware Autoupdate::
 Selects firmware image information for the Service Processors whose
automatic update matches the specified configuration.

``[-last-update-status {failed|passed}]`` - Last Update Status::
Selects firmware image information for the Service Processors whose last
update is of the specified status.

== Examples

The following command displays basic firmware information for the Service
Processors of all the nodes.
```

cluster1::> system service-processor image show
Is
Node Type Image Status Current Version
--- --- --- --- --- ---
node1 SP
primary installed true 2.2X8
backup installed false 2.2X5
node2 SP
primary installed true 2.2X8
backup installed false 2.2X5
4 entries were displayed.

cluster1::>

```
The following command displays all available firmware information for the
Service Processors of all the nodes.
```

cluster1::> system service-processor image show -instance
Node: node1
Image: primary
Type: SP
Image Status: installed
Is Image Current: true

Firmware Version: 2.2X8
Firmware Autoupdate: true
Last Update Status: passed
Node: node1
Image: backup
Type: SP
Image Status: installed
Is Image Current: false
Firmware Version: 2.2X5
Firmware Autoupdate: true
Last Update Status: passed
Node: node2
Image: primary
Type: SP
Image Status: installed
Is Image Current: true
Firmware Version: 2.2X8
Firmware Autoupdate: true
Last Update Status: passed
Node: node2
Image: backup
Type: SP
Image Status: installed
Is Image Current: false
Firmware Version: 2.2X5
Firmware Autoupdate: true
Last Update Status: passed
4 entries were displayed.

cluster1::>

```
[[ID65f145b878a986badacc854e2e0dabe9]]
= system service-processor image update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Update Service Processor firmware

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor image update` command installs a new
firmware version on the Service Processor (SP) or Baseboard Management
Controller (BMC) of specified node in a cluster. If this command fails, it
will display and log an appropriate error message and abort. No automatic
```

```
command retries will be performed. This command also specifies which
firmware image is to be installed on the SP or BMC and how.

You can use the command xref:{relative_path}system-service-processor-
image-update-progress-show.html[system service-processor image update-
progress show] to check the progress of the update.

The following parameter combinations are not supported for this command:

* `-baseline```_true_`` with `-package```_<text>_``



== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node whose SP or BMC firmware is to be
updated.

``[-package <text>]`` - Firmware Package::
This parameter specifies the package that will be installed. You can find
the package file in the SP Update Repository field of the
xref:{relative_path}system-node-image-package-show.html[system node image
package show] command. If you do not specify this parameter, the SP or BMC
is updated to the most recent version of the firmware that is available in
the update repository. You must specify this parameter if `baseline` is
false or omitted.

``[-baseline {true|false}]`` - Install Baseline::
If you set this parameter to true, the command installs the SP or BMC
firmware version that is bundled with the currently running release of
ONTAP. This is a safety mechanism that allows you to revert the SP or BMC
firmware to the version that was qualified and bundled with the currently
running version of ONTAP on your system. If not specified, this parameter
defaults to false.

``[-pre-check {true|false}]`` - Perform Pre-checks [.advaced-
priv]#(privilege: advanced)#::
If you set this parameter to ``_false_`` , the command skips pre-requisite
checks, such as target versions below baseline or compatibility tokens.
The default value for this parameter is ``_true_`` .

== Examples

The following command reverts the firmware on the SP or BMC of the local
node to the version that was packaged with the currently running release
of ONTAP. The second command displays the status of the in-progress
```

```
firmware install.
```

cluster1::> system service-processor image update -node local -baseline true

```

```

cluster1::>

cluster1::> system service-processor image update-progress show
In Percent
Node Progress Start Time Done End Time
---------- -------- ---------- ------- ----------
node1 yes 8/28/2012 20:00:34 99 -
node2 no - 0 -
2 entries were displayed.

cluster1::>

```
== Related Links

* xref:{relative_path}system-service-processor-image-update-progress-show.html[system service-processor image update-progress show]
* xref:{relative_path}system-node-image-package-show.html[system node image package show]




[[ID8a3e6ca3a8918b7a0e828bbf691607da]]
= system service-processor image update-progress show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status for the latest Service Processor firmware update

*Availability:* This command is available to _cluster_ administrators at the _admin_ privilege level.

== Description

The `system service-processor image update-progress show` command displays the progress information of firmware updates on the Service Processor (SP) or Baseboard Management Controller (BMC) of the specified nodes. The `"in-progress"` field displays "``_no_`` " if no update is in progress. This
```

```
command does not display the progress of an SP/BMC firmware update that is
triggered from the SP CLI.

== 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>|local}]`` - Node::
 This parameter displays the status of SP or BMC firmware update for the
specified node.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Latest SP Firmware Update Start
Timestamp::
This parameter displays the status of the SP or BMC whose firmware update
start time matches the specified value.

``[-percent-done <integer>]`` - Latest SP Firmware Update Percentage
Done::
This parameter displays the status of the SP or BMC whose update
completion percentage matches the specified value.

``[-end-time <MM/DD/YYYY HH:MM:SS>]`` - Latest SP Firmware Update End
Timestamp::
This parameter displays the status of the SP or BMC whose firmware update
end time matches the specified value.

``[-in-progress {yes|no}]`` - Is Update in Progress::
This parameter displays the update status of the SP or BMC that matches
the specified in-progress status.

== Examples

The following example starts a firmware update on the local node and then
uses the command `system service-processor image update-progress show` to
display progress of firmware updates on SPs or BMCs of all nodes in the
system.
```

cluster1::> system service-processor image update -node local -baseline true

cluster1::>

cluster1::> system node service-processor image update-progress show
In Percent
Node Progress Start Time Done End Time
--- --- --- --- ---
node1 yes 8/28/2012 20:00:34 99 -
node2 no - 0 -
2 entries were displayed.

cluster1::>

```
[[ID64ac17eb8936341da9cc2debfede1162]]
= system service-processor log show-allocations
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Service Processor log allocation map

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor log show-allocations` command displays the
allocation map of the Service Processor logs collected in the cluster. The
Service Processor logs of a node are archived in the mroot directory of
the collecting node. This command displays the sequence numbers for the
Service Processor log files that reside in each collecting node.

== 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>|local}]`` - Node::
If you specify this parameter, the command displays the sequence numbers
of Service Processor log files that the specified node has collected.
```

```
``[-remote-node <text>]`` - Remote Node::
If you specify this parameter, the command displays the sequence numbers
of Service Processor log files that have been collected from the specified
remote node.

``[-seqList <integer>,...]`` - Log File Sequence Numbers::
If you specify this parameter, the command displays information about
Service Processor log files with the specified sequence number.

== Examples

The following example displays the allocation map of the Service Processor
log files in the cluster.
```

cluster1::> system service-processor log show-allocation
Node From Which Node Log File Sequence
--- --- ---
cluster1-01
cluster1-01 10, 11, 12, 13, 15
cluster1-02 14, 15, 16, 17
cluster1-02
cluster1-01 14
cluster1-02 11, 12, 13
4 entries were displayed.

cluster1::>

```
[[ID1fe52c7385431a0eb4109ff9c97f70a5]]
= system service-processor network modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the network configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor network modify` command modifies the network
configuration of the Service Processor of specified node or nodes in a
cluster.

If the SP automatic network configuration has been enabled, the `system
```

```
service-processor network modify` command allows you to only enable or
disable the SP IPv4 or Ipv6 network interface.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node whose Service Processor's network
configuration is to be modified.

``-address-family {IPv4|IPv6}`` - Address Family::
This parameter specifies whether the IPv4 or the IPv6 configuration is to
be modified.

``[-enable {true|false}]`` - Interface Enabled::
This parameter enables or disables the underlying network interface for
the specified `address-family` . This is a mandatory parameter.

``[-dhcp {v4|none|v6}]`` - DHCP Status::
If this parameter is set to v4, the Service Processor uses network
configuration from the DHCP server. Otherwise, the Service Processor uses
the network address you specify. If this parameter is not set to v4 or is
not specified, you must specify the IP address, netmask, prefix-length,
and gateway in the command. DHCP is not supported for IPv6 configuration.

``[-ip-address <IP Address>]`` - IP Address::
This parameter specifies the public IP address for the Service Processor.
You must specify this parameter when the `-dhcp` parameter is not set to
v4.

{ ``[-netmask <IP Address>]`` - Netmask::
This parameter specifies the netmask for a Service Processor that uses an
IPv4 address. This parameter has no effect if the IP address family is set
to IPv6. You must specify this parameter when DHCP is not v4 and the
address family is IPv4.

| ``[-prefix-length <integer>]`` - Prefix Length of Subnet Mask }::
This parameter specifies the network prefix-length of the Service
Processor if the address family is set to IPv6. The parameter has no
effect when the address family is set to IPv4. You must specify this
parameter when DHCP is not set to v4 and when the address family is set to
IPv6.

``[-gateway <IP Address>]`` - Gateway IP Address::
This parameter specifies network gateway of the Service Processor. You
must specify this parameter when DHCP is not set to v4.

== Examples
```

```
The following example enables the network interface for IPv4 on the
Service Processor of the local node. It first displays the current network
configuration information of the local node to show the network interface
is initially disabled, and then enables it with IP address 192.168.1.202,
netmask as 255.255.255.0 and gateway as 192.168.1.1. It displays the
interim state with SP Network Setup Status field showing "in-progress". It
finally displays the network configuration again to confirm the specified
values took effect.
```

```
cluster1::> system service-processor network show -instance -node local
Node: node2
Address Family: IPv4
Interface Enabled: false
Type of Device: SP
Status: online
Link Status: disabled
DHCP Status: -
IP Address: -
MAC Address: ab:cd:ef:fe:ed:02
Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: -
Time Last Updated: Fri Jun 13 16:29:55 GMT 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: succeeded
SP Network Setup Failure Reason: -
Node: node2
Address Family: IPv6
Interface Enabled: false
Type of Device: SP
Status: online
Link Status: disabled
DHCP Status: none
IP Address: -
MAC Address: ab:cd:ef:fe:ed:02
Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: -
Time Last Updated: Fri Jun 13 16:29:55 GMT 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: not-setup
SP Network Setup Failure Reason: -
2 entries were displayed.

cluster1::>
```

```
cluster1::> system service-processor network modify -node local -address-family IPv4 -enable true -ip-address
192.168.1.202 -netmask 255.255.255.0 -gateway 192.168.1.1
```

```
cluster1::>
cluster1::> system service-processor network show -instance -node local
Node: node2
Address Family: IPv4
Interface Enabled: false
Type of Device: SP
Status: online
Link Status: disabled
DHCP Status: -
IP Address: -
MAC Address: ab:cd:ef:fe:ed:02
Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: -
Time Last Updated: Fri Jun 13 16:29:55 GMT 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: in-progress
SP Network Setup Failure Reason: -
Node: node2
Address Family: IPv6
Interface Enabled: false
Type of Device: SP
Status: online
Link Status: disabled
DHCP Status: none
IP Address: -
MAC Address: ab:cd:ef:fe:ed:02
cluster1::> system service-processor network show -instance -node local
Node: node2
Address Family: IPv4
Interface Enabled: true
Type of Device: SP
Status: online
Link Status: up
DHCP Status: none
IP Address: 192.168.1.202
MAC Address: ab:cd:ef:fe:ed:02
Netmask: 255.255.255.0
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: 192.168.1.1
Time Last Updated: Fri Jun 13 16:29:55 GMT 2014
```

Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: succeeded
SP Network Setup Failure Reason: -
Node: node2
Address Family: IPv6
Interface Enabled: false
Type of Device: SP
Status: online
Link Status: disabled
DHCP Status: none
IP Address: -
MAC Address: ab:cd:ef:fe:ed:02
Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: -
Time Last Updated: Fri Jun 13 16:29:55 GMT 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: not-setup
SP Network Setup Failure Reason: -
2 entries were displayed.

cluster1::>

```
[[ID78d0f3a002d34f70d093a56a60d1d270]]
= system service-processor network show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the network configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor network show` command displays the network
configuration of the Service Processor of each node in a cluster. You can
limit output to specific types of information and specific nodes in the
cluster, or filter output by specific field values.

In case a node is offline or its Service Processor management daemon is
down, the command displays the last known IP address of its Service
Processor. Only the IP address is displayed in such cases.
```

```
== 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>|local}]`` - Node::
Selects network configuration information for the Service Processor of the
specified node.

``[-address-family {IPv4|IPv6}]`` - Address Family::
Selects network configuration information for the Service Processors that
have the specified IP address family.

``[-enable {true|false}]`` - Interface Enabled::
Selects network configuration information for the Service Processors whose
network interface for the given `address-family` is enabled or disabled as
specified.

``[-type {SP|NONE|BMC}]`` - Type of Device::
Selects network configuration information for the Service Processors of
the specified type.

``[-status {online|offline|sp-daemon-offline|node-
offline|degraded|rebooting|unknown|updating}]`` - Status::
Selects network configuration information for the Service Processors whose
status matches the specified value.

``[-link-status {up|down|disabled|unknown}]`` - Link Status::
Selects network configuration information for the Service Processors whose
link status matches the specified value.

``[-dhcp {v4|none|v6}]`` - DHCP Status::
Selects network configuration information for the Service Processors whose
DHCP status matches the specified value.

``[-ip-address <IP Address>]`` - IP Address::
Selects network configuration information for the Service Processors that
use the specified IP address.

``[-mac <MAC Address>]`` - MAC Address::
```

```
Selects network configuration information for the Service Processors that
use the specified MAC address.

``[-netmask <IP Address>]`` - Netmask::
This parameter displays information only for the Service Processors that
use the specified netmask.

``[-prefix-length <integer>]`` - Prefix Length of Subnet Mask::
Selects network configuration information for the Service Processors whose
prefix length of subnet mask matches the specified value.

``[-router-ip <IP Address>]`` - Router Assigned IP Address::
Selects network configuration information for the Service Processors whose
router-assigned IP address matches the specified value.

``[-link-local-ip <IP Address>]`` - Link Local IP Address::
Selects network configuration information for the Service Processors whose
link local IP address matches the specified value.

``[-gateway <IP Address>]`` - Gateway IP Address::
Selects network configuration information for the Service Processors whose
gateway IP address matches the specified value.

``[-time-last-updated <text>]`` - Time Last Updated::
Selects network information for the Service Processors that have the
specified time stamp showing when configuration was last updated.

``[-subnet-name <text>]`` - Subnet Name::
Selects network information for the Service Processors that use the
specified subnet-name for SP automatic configuration.

``[-is-ipv6-ra-enabled {true|false}]`` - Enable IPv6 Router Assigned
Address::
Selects network information for the Service Processors that have the
specified status for IPv6 router-assigned address.

``[-setup-status {not-setup|succeeded|in-progress|failed}]`` - SP Network
Setup Status::
Selects network information for the Service Processors that have the
specified status for network interface setup.

``[-setup-failure-reason {success|subnet-out-of-address|invalid-
subnet|other-error}]`` - SP Network Setup Failure Reason::
Selects network information for the Service Processors that have the
specified reason for network interface setup failure.

== Examples
```

```
The following example displays basic network configuration information for
the Service Processors of all the nodes.
```

```
cluster1::> system service-processor network show
Address
Node Status Type Link State IP Address
------------- -------------- --------- ----------- ------------------------
node1 online IPv4 up 192.168.1.201
DHCP: v4
MAC Address: ab:cd:ef:fe:ed:01
Network Gateway: 192.168.1.1
Network Mask (IPv4 only): 255.255.255.0
Prefix Length (IPv6 only): -
IPv6 RA Enabled: -
Subnet Name: -
SP Network Setup Status: succeeded
node1 online IPv6 disabled -
DHCP: none
MAC Address: ab:cd:ef:fe:ed:01
Network Gateway: -
Network Mask (IPv4 only): -
Prefix Length (IPv6 only): -
IPv6 RA Enabled: -
Subnet Name: -
SP Network Setup Status: not-setup
node2 online IPv4 up 192.168.1.202
DHCP: v4
MAC Address: ab:cd:ef:fe:ed:02
Network Gateway: 192.168.1.1
Network Mask (IPv4 only): 255.255.255.0
Prefix Length (IPv6 only): -
IPv6 RA Enabled: -
Subnet Name: -
SP Network Setup Status: succeeded
node2 online IPv6 disabled -
DHCP: none
MAC Address: ab:cd:ef:fe:ed:02
Network Gateway: -
Network Mask (IPv4 only): -
Prefix Length (IPv6 only): -
IPv6 RA Enabled: -
Subnet Name: -
SP Network Setup Status: not-setup
4 entries were displayed.

cluster1::>
```

```
The following example displays all available network configuration
information for the Service Processors of all the nodes.
```

```
cluster1::> system service-processor network show -instance
Node: node1
Address Family: IPv4
Interface Enabled: true
Type of Device: SP
Status: online
Link Status: up
DHCP Status: v4
IP Address: 192.168.1.201
MAC Address: ab:cd:ef:fe:ed:01
Netmask: 255.255.255.0
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: 192.168.1.1
Time Last Updated: Fri Jun 13 17:03:59 GMT 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: succeeded
SP Network Setup Failure Reason: -
Node: node1
Address Family: IPv6
Interface Enabled: false
Type of Device: SP
Status: online
Link Status: disabled
DHCP Status: none
IP Address: -
MAC Address: ab:cd:ef:fe:ed:01
Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: -
Time Last Updated: Fri Jun 13 17:03:59 GMT 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: not-setup
SP Network Setup Failure Reason: -
Node: node2
Address Family: IPv4
Interface Enabled: true
Type of Device: SP
Status: online
Link Status: up
DHCP Status: v4
IP Address: 192.168.1.202
MAC Address: ab:cd:ef:fe:ed:02
Netmask: 255.255.255.0
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: 192.168.1.1
Time Last Updated: Fri Jun 13 17:03:59 GMT 2014
```

```
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: succeeded
SP Network Setup Failure Reason: -
Node: node2
Address Family: IPv6
Interface Enabled: false
Type of Device: SP
Status: online
Link Status: disabled
DHCP Status: none
IP Address: -
MAC Address: ab:cd:ef:fe:ed:02
Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: -
Time Last Updated: Fri Jun 13 17:03:59 GMT 2014
Subnet Name: -
Enable IPv6 Router Assigned Address: -
SP Network Setup Status: not-setup
SP Network Setup Failure Reason: -
4 entries were displayed.

cluster1::>
```

```
[[ID2a357ef654da747eec31613cce426ebe]]
= system service-processor network auto-configuration disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable Service Processor Auto-Configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor network auto-configuration disable` command
disables the SP's use of subnet resource for the automatic configuration
of its networking port. This command is a cluster-wide configuration. When
you disable the SP automatic network configuration, all SPs in the cluster
will be configured to use IPv4 DHCP. Any addresses previously allocated
from the subnet to the SP will be released. If the SP fails to obtain an
IPv4 IP address from the DHCP server, an EMS message warns you about the
failure. The IPv6 interface will be disabled.

== Parameters

``-address-family {IPv4|IPv6}`` - Subnet Address Family::
This parameter specifies whether the IPv4 or the IPv6 automatic
configuration is to be disabled for the SP.

== Examples

The following example disables the automatic configuration for IPv4 on the
SP. It first displays the current network configuration and then disables
the SP IPv4 automatic network configuration.
```

cluster1::>system service-processor network show
Address
Node Status Family Link State IP Address
---------- ---------- --------- ---------- ----------
node1
online IPv4 up 192.168.1.2
DHCP: none
MAC Address: ab:cd:ef:fe:ed:01
Network Gateway: 192.168.1.1
Network Mask (IPv4 only): 255.255.255.0
Prefix Length (IPv6 only): -

```
IPv6 RA Enabled: -
Subnet Name: ipv4_test
SP Network Setup Status: succeeded
```

```
cluster1::>system service-processor network auto-configuration disable -address-family Ipv4

cluster1::>system service-processor network auto-configuration show
Cluster Name SP IPv4 Subnet Name SP IPv6 Subnet Name
-------------------- ---------------------------- ----------------------------
cluster1 - -
```

```
cluster1::>system service-processor network show
Address
Node Status Family Link State IP Address
------------- -------------- --------- ----------- ------------------------
node1
online IPv4 up 192.168.1.184
DHCP: v4
MAC Address: ab:cd:ef:fe:ed:01
Network Gateway: 192.168.1.1
Network Mask (IPv4 only): 255.255.255.0
Prefix Length (IPv6 only): -
IPv6 RA Enabled: -
Subnet Name: -
SP Network Setup Status: succeeded
```

```
[[IDbe3c57fe44a276678d81c53d2cb908a0]]
= system service-processor network auto-configuration enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable Service Processor Auto-Configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor network auto-configuration enable` command
enables the automatic network configuration for the SP. This is a cluster-
wide configuration. Every node in the cluster will use the specified
subnet to allocate IP address, subnet mask and gateway address for the SP
configuration. When the SP automatic network configuration is enabled, you
do not need to manually manage the SP network of individual nodes. A node
that subsequently joins the cluster uses the specified subnet to configure
its SP network automatically.

Prior to running this command, the subnet you want to use for the SP
automatic network configuration must already be defined in the cluster and
must have no resource conflicts with the SP network interface.

== Parameters

``-address-family {IPv4|IPv6}`` - Subnet Address Family::
This parameter specifies whether the IPv4 or the IPv6 automatic
configuration is to be enabled for the SP.

``-subnet-name <text>`` - Subnet Name::
This parameter specifies the network subnet that the SP will use for
automatic network configuration.

== Examples

The following example enables the automatic network configuration for IPv4
on the SP. It first displays the current SP network configuration,
displays available network subnet in the cluster, and then enable the SP
to use the subnet for IPv4 automatic configuration.
```

```
cluster1::>system service-processor network show
Address
Node Status Family Link State IP Address
------------- -------------- --------- ----------- ------------------------
node1
online IPv4 up 192.168.1.201
DHCP: v4
MAC Address: ab:cd:ef:fe:ed:01
Network Gateway: 192.168.1.1
Network Mask (IPv4 only): 255.255.255.0
Prefix Length (IPv6 only): -
IPv6 RA Enabled: -
Subnet Name: -
SP Network Setup Status: succeeded
```

```
cluster1::> network subnet show

IPspace: Default
Subnet Broadcast Avail/
Name Subnet Domain Gateway Total Ranges
--------- ---------------- --------- --------------- --------- ---------------
ipv4_test 192.168.1.0/24 Default 192.168.1.1 3/5 192.168.1.2-192.168.1.6
```

```
cluster1::>system service-processor network auto-configuration enable -address-family ipv4 -subnet-name
ipv4_test

cluster1::system service-processor network> show
Address
Node Status Family Link State IP Address
------------- -------------- --------- ----------- ------------------------
node1
online IPv4 up 192.168.1.2
DHCP: none
MAC Address: ab:cd:ef:fe:ed:01
Network Gateway: 192.168.1.1
Network Mask (IPv4 only): 255.255.255.0
Prefix Length (IPv6 only): -
IPv6 RA Enabled: -
Subnet Name: ipv4_test
SP Network Setup Status: succeeded
```

```
[[IDe5e9d97e65ab562407d69038e0834e4d]]
= system service-processor network auto-configuration show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Service Processor Auto-Configuration Setup

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor network auto-configuration show` command
displays the names of the IPv4 and IPv6 network subnet objects configured
in the cluster that the SP uses for automatic configuration.

== Examples

The following example shows that the SP is configured to use the
"ipv4_test" IPv4 subnet in the cluster for the SP automatic network
configuration.
```

cluster1::>system service-processor network auto-configuration show
Cluster Name SP IPv4 Subnet Name SP IPv6 Subnet Name
--- --- ---
cluster1 ipv4_test -

```
[[IDf4bbc52833d5b80cd2655a97026ed3e7]]
= system service-processor ssh add-allowed-addresses
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add IP addresses to the list that is allowed to access the Service
Processor

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor ssh add-allowed-addresses` command grants IP
addresses access to the Service Processor.

== Parameters

``-allowed-addresses <IP Address/Mask>,...`` - Public IP Addresses::
Use this parameter to specify one or more IP addresses with corresponding
netmasks. The value should be specified in the format of address/netmask,
for example, 10.98.150.10/24, fd20:8b1e:b255:c09b::/64. Use commas to
separate multiple address/netmask pairs. If "0.0.0.0/0, ::/0" is specified
in the parameter, any IP address is allowed to access the Service
Processor.

== Examples

The following examples grant the specified IP addresses access to the
Service Processor and display the list of public IP addresses that are
allowed to access the Service Processor.
```

cluster1::> system service-processor ssh show
Allowed Addresses: 0.0.0.0/0, ::/0

cluster1::> system service-processor ssh add-allowed-addresses -allowed-addresses 192.168.1.202/24, 192.168.10.201/24
Warning: The default "allow all" setting (0.0.0.0/0, ::/0) will be replaced
with your changes. Do you want to continue? {y|n}: y

cluster1::> system service-processor ssh show
Allowed Addresses: 192.168.1.202/24, 192.168.10.201/24

The following example enables all IP addresses to access the Service Processor.

cluster1::> system service-processor ssh add-allowed-addresses -allowed-addresses 0.0.0.0/0, ::/0

cluster1::> system service-processor ssh show
Allowed Addresses: 0.0.0.0/0, ::/0
cluster1::>

```
[[ID46a10c50d607be5de7e76281ed2791f2]]
= system service-processor ssh remove-allowed-addresses
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove IP addresses from the list that is allowed to access the Service
Processor

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor ssh remove-allowed-addresses` command blocks
the specified IP address from accessing the Service Processor. If all IP
addresses are removed from the access list, then the Service Processor is
not accessible from any IP address.

== Parameters

``-allowed-addresses <IP Address/Mask>,...`` - Public IP Addresses::
Use this parameter to specify one or more IP addresses with corresponding
netmasks. The value should be specified in the format of address/netmask,
for example, 10.98.150.10/24, fd20:8b1e:b255:c09b::/64. Use commas to
separate multiple address/netmask pairs.

== Examples

The following example prevents the specified IP addresses from accessing
the Service Processor. It also displays the list of public IP addresses
that are allowed to access the Service Processor.
```

cluster1::> system service-processor ssh show
Allowed Addresses: 192.168.1.202/24, 192.168.10.201/24

cluster1::> system service-processor ssh remove-allowed-addresses -allowed-addresses 192.168.1.202/24, 192.168.10.201/24

Warning: If all IP addresses are removed from the allowed address list, all IP addresses will be denied access. To restore the "allow all" default, use the "system service-processor ssh add-allowed-addresses -allowed-addresses 0.0.0.0/0, ::/0" command. Do you want to continue? {y|n}: y

cluster1::> system service-processor ssh show
Allowed Addresses: -

cluster1::>

```
[[IDf8773c5b1db23730b52845fbe797c808]]
= system service-processor ssh show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display SSH security information about the Service Processor

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system service-processor ssh show` command displays the IP addresses
that are allowed to access the Service Processor by using SSH.

== Examples

The following example displays SSH security information about the Service
Processor.
```

cluster1::> system service-processor ssh show
Allowed Addresses: 0.0.0.0/0, ::/0

cluster1::>

```
:leveloffset: -1


= system services commands

:leveloffset: +1


[[ID7eaa1170a920c0a47fb71745e8e4b1e5]]
= system services firewall modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify firewall status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services firewall modify` command modifies a node's firewall
configuration.

== Parameters

``-node {<nodename>|local}`` - Node::
Use this parameter to specify the node on which to modify firewall
configuration.

``[-enabled {true|false}]`` - Service Enabled::
Use this parameter to specify whether firewall protection is enabled
("``_true_`` ") or disabled ("``_false_`` ") for the node's network ports.
The default setting is `true` .

== Examples

The following example enables firewall protection for a node named node1:
```

cluster1::> system services firewall modify -node node1 -enabled true

```
[[ID88cfa1e161aed4989540f2d60e0215fd]]
= system services firewall show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show firewall status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services firewall show` command displays firewall
configuration. If the command is issued without any parameters, it
displays information about all nodes in the cluster. You can also query
specific nodes for their firewall information by running the command with
the `-node` parameter.

== 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>|local}]`` - Node::
Selects information about the firewall settings on the node you specify.

``[-enabled {true|false}]`` - Service Enabled::
Selects information about the nodes with the firewall enabled ("``_true_``
") or disabled ("``_false_`` ").

== Examples

The following example displays information about firewall configuration
for all nodes in the cluster:
```

cluster1::> system services firewall show

```
Node Enabled
-------------- -------
node0 true
node1 true
node2 true
node3 true
4 entries were displayed.
```

```
[[ID8f914ddd7a0c0daef593ba5f81c3bbc9]]
= system services firewall policy clone
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Clone an existing firewall policy

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future ONTAP
release. Use xref:{relative_path}network-interface-service-policy-
clone.html[network interface service-policy clone] instead.

The `system services firewall policy clone` command creates a new firewall
policy that is an exact copy of an existing policy, but has a new name.

== Parameters

``-vserver <text>`` - Vserver owning the Policy::
Use this parameter to specify the name of the Vserver owning the existing
policy to copy.

``-policy <text>`` - Firewall Policy to be Cloned::
Use this parameter to specify the name of the existing policy to copy.

``[-destination-vserver <text>]`` - Vserver owning the New Firewall
Policy::
Use this parameter to specify the name of the Vserver that will own the
new policy to create.

``-destination-policy <text>`` - Name of New Firewall Policy::
Use this parameter to specify the name of the new policy to create.

== Examples

This example creates a new firewall policy named "data2" on Vserver "vs0"
from an existing firewall policy named "data" on Vserver "vs1".
```

cluster1::> system services firewall policy clone -vserver vs0 -policy data -destination-vserver vs1 -destination -policy data2

```
== Related Links

* xref:{relative_path}network-interface-service-policy-clone.html[network
interface service-policy clone]




[[ID79eaf188d329c831f8bf913d18861317]]
= system services firewall policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Create a firewall policy entry for a network service

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future ONTAP
release. Use xref:{relative_path}network-interface-service-policy-
create.html[network interface service-policy create] or
xref:{relative_path}network-interface-service-policy-add-
service.html[network interface service-policy add-service] instead.

The `system services firewall policy create` command creates a firewall
policy entry with the specified name and network service. This command is
used both to create the first network service associated with a new
firewall policy, and to add to an existing firewall policy by associating
another network service with an existing policy. You can optionally
specify one or more IP addresses with corresponding netmasks that are
allowed to use the firewall policy entry.

You can use the xref:{relative_path}network-interface-modify.html[network
interface modify] command with the `-firewall-policy` parameter to put a
firewall policy into effect for a given logical interface by modifying
that logical interface to use the specified firewall policy.
```

```
NOTE: You can use this command to create an empty firewall policy by
creating a single policy entry for the "none" firewall service. When used
by a logical network interface (LIF), an empty firewall policy will block
all services managed using firewall policies.



== Parameters

``-vserver <vserver>`` - Vserver Name::
Use this parameter to specify the name of the Vserver on which the policy
is to be created.

``-policy <textpolicy_name>`` - Policy::
Use this parameter to specify the name of the policy that is to be
created.

``-service <service>`` - Service::
Use this parameter to specify the network service that is associated with
the policy. Possible values include:
+
  * dns - The DNS protocol server
  * http - The HTTP protocol
  * ndmp - The NDMP tape backup protocol
  * ndmps - The NDMPS tape backup protocol
  * none - No protocol (for creating an empty policy)
  * ntp - The NTP protocol
  * rsh - The RSH protocol
  * snmp - The SNMP protocol
  * telnet - The Telnet protocol



``-allow-list <IP Address/Mask>,...`` - Allowed IPs::
Use this parameter to specify one or more IP addresses with corresponding
netmasks that are to be allowed by this firewall policy. The correct
format for this parameter is address/netmask, similar to "192.0.2.128/25".
Multiple address/netmask pairs should be separated with commas. Use the
value `0.0.0.0/0` for "any".

== Examples

The following example creates a firewall policy named data that uses the
NDMP protocol and enables access from all IP addresses on the
192.0.2.128/25 subnet:
```

cluster1::> system services firewall policy create -policy data -service ndmp -allow-list 192.0.2.128/25

```
The following example adds an entry to the firewall policy named data,
associating the DNS protocol with that policy and enabling access from all
IP addresses on the 192.0.2.128/25 subnet:
```

cluster1::> system services firewall policy create -policy data -service dns -allow-list 192.0.2.128/25

```
== Related Links

* xref:{relative_path}network-interface-service-policy-create.html[network
interface service-policy create]
* xref:{relative_path}network-interface-service-policy-add-
service.html[network interface service-policy add-service]
* xref:{relative_path}network-interface-modify.html[network interface
modify]




[[IDd9a59150fd5f8bf51a6cc275e4ad27ed]]
= system services firewall policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Remove a service from a firewall policy

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future ONTAP
release. Use xref:{relative_path}network-interface-service-policy-
delete.html[network interface service-policy delete] or
xref:{relative_path}network-interface-service-policy-remove-
service.html[network interface service-policy remove-service] instead.

The `system services firewall policy delete` command deletes a firewall
policy. You cannot delete a policy that is being used by a logical
interface. Use the xref:{relative_path}network-interface-
```

```
modify.html[network interface modify] command with the `-firewall-policy`
parameter to change a network interface's firewall policy.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Use this parameter to specify the Vserver of the policy to delete.

``-policy <textpolicy_name>`` - Policy::
Use this parameter to specify the name of the policy to delete.

``-service <service>`` - Service::
Use this parameter to specify the policy's network service to delete.

== Examples

The following example deletes a firewall policy that uses the Telnet
protocol on the policy named data:
```

cluster1::> system services firewall policy delete -policy data -service telnet

```
Use wildcards to delete entire policies at once, or particular services
from every policy. This example deletes the entire intercluster policy.
```

cluster1::> system services firewall policy delete -policy intercluster -service *

```
This example deletes the telnet service from every policy.
```

cluster1::> system services firewall policy delete -policy * -service telnet

```
== Related Links

* xref:{relative_path}network-interface-service-policy-delete.html[network
interface service-policy delete]
* xref:{relative_path}network-interface-service-policy-remove-
service.html[network interface service-policy remove-service]
* xref:{relative_path}network-interface-modify.html[network interface
modify]




[[ID0b1c307a24a23abcf02e8ee926a5269f]]
= system services firewall policy modify
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Modify a firewall policy entry for a network service

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future ONTAP
release. Use xref:{relative_path}network-interface-service-policy-modify-
service.html[network interface service-policy modify-service] instead.

The xref:{relative_path}system-services-firewall-modify.html[system
services firewall modify] command enables you to modify the list of IP
addresses and netmasks associated with a firewall policy.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Use this parameter to specify the Vserver of the policy to modify.

``-policy <textpolicy_name>`` - Policy::
Use this parameter to specify the name of the policy to modify.

``-service <service>`` - Service::
Use this parameter to specify the policy's network service to modify.

``[-allow-list <IP Address/Mask>,...]`` - Allowed IPs::
Use this parameter to specify one or more IP addresses with corresponding
netmasks that are allowed by this firewall policy. The correct format for
this parameter is address/netmask, similar to "192.0.2.128/25". Multiple
address/netmask pairs should be separated with commas. Use the value
`0.0.0.0/0` for "any".

== Examples

The following example modifies the firewall policy named data that uses
the NDMP protocol to enable access from all addresses on the 192.0.2.128
subnet:
```

cluster1::> system services firewall policy modify -policy data -service ndmp -allow-list 192.0.2.128/25

```
== Related Links

* xref:{relative_path}network-interface-service-policy-modify-
service.html[network interface service-policy modify-service]
* xref:{relative_path}system-services-firewall-modify.html[system services
firewall modify]




[[IDf9bf4ffb50d20e89b8c39415a19df150]]
= system services firewall policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Show firewall policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future ONTAP
release. Use xref:{relative_path}network-interface-service-policy-
show.html[network interface service-policy show] instead.

The `system services firewall policy show` command displays information
about firewall policies.

NOTE: Some firewall policies contain a single entry for the "none"
firewall service. You can consider these policies to be empty. When used
by a logical network interface (LIF), an empty firewall policy will block
all services managed using firewall policies.



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
```

```
displays only the fields that you specify.

| ``[-instance  ]`` }::
Use this parameter to display all the fields for the specified policies.

``[-vserver <vserver>]`` - Vserver Name::
Use this parameter to display information only about the Vserver you
specify.

``[-policy <textpolicy_name>]`` - Policy::
Use this parameter to display information about the policy you specify.

``[-service <service>]`` - Service::
Use this parameter to display information about the services you specify.

``[-allow-list <IP Address/Mask>,...]`` - Allowed IPs::
Use this parameter to display information about the firewall policies that
match the list of allowed IP addresses and netmasks you specify. The
correct format for this parameter is address/netmask, similar to
"192.0.2.128/25". Multiple address/netmask pairs should be separated with
commas.

``[-ipspace <text>]`` - IPspace::
Use this parameter to display information only about the IPspace you
specify.

== Examples

The following example displays information about all firewall policies:
```

```
cluster1::> system services firewall policy show
Vserver Policy Service Allowed
------- ------------ ---------- -------------------
cluster1
data
dns 0.0.0.0/0, ::/0
ndmp 0.0.0.0/0, ::/0
ndmps 0.0.0.0/0, ::/0
cluster1
intercluster
ndmp 0.0.0.0/0, ::/0
ndmps 0.0.0.0/0, ::/0
cluster1
mgmt
dns 0.0.0.0/0, ::/0
http 0.0.0.0/0, ::/0
ndmp 0.0.0.0/0, ::/0
ndmps 0.0.0.0/0, ::/0
ntp 0.0.0.0/0, ::/0
```

snmp 0.0.0.0/0, ::/0
cluster1
mgmt-nfs
dns 0.0.0.0/0, ::/0
http 0.0.0.0/0, ::/0
ndmp 0.0.0.0/0, ::/0
ndmps 0.0.0.0/0, ::/0
ntp 0.0.0.0/0, ::/0
snmp 0.0.0.0/0, ::/0
17 entries were displayed.

cluster1::>

```
== Related Links

* xref:{relative_path}network-interface-service-policy-show.html[network
interface service-policy show]




[[IDa5a5f2b56ec813125660bdea403052e5]]
= system services manager install show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of installed services

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services manager install show` command displays information
about installed services.

== 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.

``[-service <text>]`` - Service::
Selects information about installed services that have the name you
specify.

``[-version <service version>]`` - Version::
Selects information about installed services that have the version number
you specify.

``[-constituent <text>]`` - Constituent::
Selects information about installed services that have the constituent
process you specify.

``[-nodes {<nodename>|local}]`` - Nodes::
Selects information about services that are installed on the nodes you
specify.

``[-description <text>]`` - Description::
Selects information about installed services that match the description
you specify.

== Examples

The following example shows typical output from a two-node cluster.
```

cluster1::> system services manager install show
Service Version Constituent Nodes
---------- ------- ---------- ----------
diagnosis
1.0 schmd node1, node2
1.0 shmd node1, node2
2 entries were displayed.

```
[[IDe4cf3e878265731f9f3b3319b6af4ec2]]
= system services manager policy add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a new service policy

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services manager policy add` command adds a new service policy
to the services manager. Policies determine which versions of a service
can run on the nodes of the cluster.

== Parameters

``-service <text>`` - Service::
Use this parameter to specify the name of the service for which to add a
policy.

``-version <service version>`` - Version::
Use this parameter to specify the minimum version number of the service to
run.

== Examples

This example adds a service manager policy for version 1.0 of the
diagnosis service.
```

cluster1::> system services manager policy add -service diagnosis -version 1.0

```
[[ID4fd52eff096d85db592cd171c69c712a]]
= system services manager policy remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a service policy

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services manager policy remove` command removes a policy from
the services manager. Policies determine which versions of a service can
run on the nodes of the cluster.

== Parameters

``-service <text>`` - Service::
Use this parameter to specify the name of the service from which to remove
a policy.

``-version <service version>`` - Version::
Use this parameter to specify the version number that is configured by the
policy to remove.

== Examples

The following example shows the removal of the service policy for version
1.0 of the diagnosis service.
```

cluster1::> system services manager policy remove -service diagnosis -version 1.0

```
[[ID383601af9ad82e62fa6869a82643be1b]]
= system services manager policy setstate
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable/disable a service policy

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services manager policy setstate` command enables or disables
services manager policies. Use the xref:{relative_path}system-services-
manager-policy-show.html[system services manager policy show] command to
display information about configured policies.

== Parameters

``-service <text>`` - Service::
Use this parameter to set the state of the policy you specify.

``-version <service version>`` - Version::
Use this parameter to set the state of the policy with the version number
you specify.

``-state {on|off}`` - State::
Use this parameter with the value "on" to enable the policy. Use this
parameter with the value "off" to disable the policy.

== Examples

The following example sets the policy for version 1.0 of the diagnosis
service to off.
```

cluster1::> system services manager policy setstate -service diagnosis -version 1.0 -state off

```
== Related Links

* xref:{relative_path}system-services-manager-policy-show.html[system
services manager policy show]
```

```
[[ID63e20e619669d8352302243311845c5e]]
= system services manager policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display service policies

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services manager policy show` command displays information
about policies that determine which versions of a service can run on the
nodes of the cluster.

Use the xref:{relative_path}system-services-manager-status-
show.html[system services manager status show] command to view information
about services that are configured to run in the cluster.

== 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.

``[-service <text>]`` - Service::
Selects policies that apply to the service you specify.

``[-version <service version>]`` - Version::
Selects policies that have the version number you specify.

``[-constituent <text>]`` - Constituent::
Selects policies that have the constituent process you specify.

``[-state {on|off}]`` - State::
```

```
Use this parameter with the value "on" to select information about
policies that are currently active. Use this parameter with the value
"off" to select information about policies that are not currently active.

``[-num-active <integer>]`` - Number Active::
Selects policies that have the number of active (running) instances you
specify.

``[-target-nodes <service affinity>,...]`` - Target Nodes::
Selects policies that are configured to run on the nodes you specify.

``[-tag <UUID>]`` - Tag [.advaced-priv]#(privilege: advanced)#::
Selects policies that have the UUID you specify. Use this parameter with
the `-fields` parameter to display a list of the UUIDs of configured
services.

== Examples

The following example shows typical output for this command.
```

cluster1::> system services manager policy show
Service Version State Constituent Number Target
Active Nodes
---------- ------- ----- ---------- ------ ----------
diagnosis
1.0 on schmd 1 any
1.0 on shmd 1 any
2 entries were displayed.

```
== Related Links

* xref:{relative_path}system-services-manager-status-show.html[system
services manager status show]




[[ID9a1ebfbe749436733315490d480c3a51]]
= system services manager status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the status of a service
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services manager status show` command displays the status of
system services that are configured to run in the cluster.

System services run on the nodes of the cluster based on policies.
Policies determine which versions of a service can run on the nodes of the
cluster. Use the xref:{relative_path}system-services-manager-policy-
show.html[system services manager policy show] command to view existing
policies.

== 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.

``[-service <text>]`` - Service::
Selects information about services that match the service name you
specify.

``[-version <service version>]`` - Version::
Selects information about services that are configured to run the version
number you specify. The configured version is the minimum version that is
allowed to run in the cluster according to a policy. Use the
xref:{relative_path}system-services-manager-policy-show.html[system
services manager policy show] command to view information about service
policies.

``[-constituent <text>]`` - Constituent::
Selects information about services that have the constituent process you
specify.

``[-actual-version <service version>]`` - Actual Version::
Selects information about services that are running the version number you
specify. This number can be higher than the configured version if a more
recent version is installed on the node that is running the service.

``[-node <nodename>]`` - Node::
```

```
Selects information about services that the services manager has assigned
to run on the nodes you specify. If the service state is "running", the
service is running on these nodes.

``[-state <svc_state>]`` - State::
Selects information about services that are in the state you specify.

``[-is-running {true|false}]`` - Is Running::
Use this parameter with the value "true" to select information about
services that are currently running. Use this parameter with the value
"false" to select information about services that are not currently
running.

== Examples

The example below shows typical output for a simple cluster.
```

cluster1::> system services manager status show
Service Version Constituent Actual Node State
Version
---------- ------- ---------- ------- ---------- ----------
diagnosis
1.0 schmd 1.0 cluster1-01 running
1.0 shmd 1.0 cluster1-01 running
2 entries were displayed.

```
== Related Links

* xref:{relative_path}system-services-manager-policy-show.html[system
services manager policy show]




[[ID2305d95780c2636aaa71671d2777828b]]
= system services ndmp kill-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Kill all NDMP sessions

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services ndmp kill-all` command is used to terminate all NDMP
sessions on a particular node in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node::
Node on which all NDMP sessions needs to be terminated.

== Examples

The following example shows how all NDMP sessions on the node named node1
can be terminated:
```

cluster1::> system services ndmp kill-all -node node1

```
[[ID914f3d5908a01f92d1510d492e738b36]]
= system services ndmp kill
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Kill the specified NDMP session

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services ndmp kill` command is used to terminate a specific
NDMP session on a particular node in the cluster.

== Parameters

``<integer>`` - Session Identifier::
Session ID of the NDMP session.

== Examples

The following example shows how a specific NDMP session on the node named
node1 can be terminated:
```

cluster1::> system services ndmp kill 4323 -node node1

```
[[ID0ee1bda788e1e1710c971a246a48363c]]
= system services ndmp modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Modify NDMP service configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "link:vserver-services-ndmp-modify.html[vserver services
ndmp modify] " command.

The `system services ndmp modify` command allows you to modify the NDMP
configurations for a node in the cluster. One or more of the following
configurations can be modified:

* Enable/disable NDMP service
* Enable/disable sending the NDMP password in clear text. Note that MD5
authentication mode is always enabled.
* NDMP user ID



== Parameters

``-node {<nodename>|local}`` - Node::
This specifies the node whose NDMP configuration is to be modified.

``[-enable {true|false}]`` - NDMP Service Enabled::
This optionally specifies whether NDMP is enabled on the node. The default
setting is `true` .

``[-clear-text {true|false}]`` - Allow Clear Text Password::
This optionally specifies whether the NDMP password can be sent in clear
text. The default setting is `true` .

``[-user-id <text>]`` - NDMP User ID::
This optionally specifies the ID of the NDMP user.

== Examples

The following example modifies the NDMP configuration on a node named
node1. The configuration enables NDMP, disables sending the password in
clear text, and specifies an NDMP user named ndmp:
```

cluster1::> system services ndmp modify -node node1 -enable true
-clear-text false -user-id ndmp

```
== Related Links

* xref:{relative_path}vserver-services-ndmp-modify.html[vserver services
ndmp modify]




[[IDfc32d5e3f1f094cd335ae3ab386cd998]]
= system services ndmp off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Disable NDMP service

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "link:vserver-services-ndmp-off.html[vserver services ndmp
off] " command.

The `system services ndmp off` command is used to disable the NDMP service
on any node in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node::
The specific node on which NDMP service is to be disabled.

== Examples

The following example is used to turn off the NDMP service on node named
node1:
```

cluster1::> system services ndmp off -node node1

```
== Related Links

* xref:{relative_path}vserver-services-ndmp-off.html[vserver services ndmp
off]




[[ID181167e178b3164960279e50c27e631e]]
= system services ndmp on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Enable NDMP service

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "link:vserver-services-ndmp-on.html[vserver services ndmp
on] " command.

The `system services ndmp on` command is used to enable the NDMP service
across any node in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node::
The specific node on which the NDMP service is to be enabled.

== Examples

The following example is used to turn on the NDMP service on node named
node1:
```

cluster1::> system services ndmp on -node node1

```
== Related Links

* xref:{relative_path}vserver-services-ndmp-on.html[vserver services ndmp
on]




[[ID858ce533d557ee9b03dfdd29cf6b72f0]]
= system services ndmp password
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Change the NDMP password for the node

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "link:vserver-services-ndmp-generate-password.html[vserver
services ndmp generate-password] " command.

The `system services ndmp password` command is used to change the NDMP
password for a node in the cluster.

== Parameters

``-node {<nodename>|local}`` - Node::
The specific node for which the password is to be changed.

== Examples

The following example is used to change the NDMP password for the node
named node1:
```

cluster1::> system services ndmp password -node node1

Please enter password:
Confirm password:

```
== Related Links

* xref:{relative_path}vserver-services-ndmp-generate-password.html[vserver
services ndmp generate-password]




[[ID4fd02b6e5b3a1797a76ee2e3f47d6a8f]]
= system services ndmp probe
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of NDMP sessions

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services ndmp probe` command displays diagnostic information
about all the NDMP sessions in the cluster. The following fields are
displayed for each of the sessions:

* Node
* Session identifier
* NDMP version
* Session authorized
* Data state
* Data operation
* Data server halt reason
* Data server connect type
* Data server connect address
* Data server connect port
* Data bytes processed
* Mover state
* Mover mode
* Mover pause reason
* Mover halt reason
* Mover record size
* Mover record number
* Mover bytes moved
* Mover seek position
```

```
* Mover bytes left to read
* Mover window offset
* Mover window length
* Mover position
* Mover SetRecordSize flag
* Mover SetWindow flag
* Mover connect type
* Mover connect address
* Mover connect port
* Effective host
* NDMP client address
* NDMP client port
* SCSI device ID
* SCSI hostadapter
* SCSI target ID
* SCSI LUN ID
* Tape device
* Tape mode
* Is Secure Control Connection
* Data Backup Mode
* Data Path
* NDMP Source Address



== Parameters

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays information about the
sessions running on the specified node only. Node should be a valid node
name.

``[-session-id <integer>]`` - Session Identifier::
If this parameter is specified, the command displays information only
about the specified session.

``[-ndmp-version <integer>]`` - NDMP Version::
This parameter refers to the NDMP protocol version being used in the
session.

``[-session-authorized {true|false}]`` - Session Authorized::
This parameter indicates whether an NDMP session is authenticated or not.

``[-data-state <component state>]`` - Data State::
This parameter identifies the current state of the data server's state
machine.
```

```
``[-data-operation <data operation>]`` - Data Operation::
This parameter identifies the data server's current operation.

``[-data-halt-reason <halt reason>]`` - Data Server Halt Reason::
This parameter identifies the event that caused the data server state
machine to enter the HALTED state.

``[-data-con-addr-type <address type>]`` - Data Server Connect Type::
This parameter specifies the type of data connection established by the
data server. The data connection can be established locally within a given
system or between remote networked systems.

``[-data-con-addr <text>]`` - Data Server Connect Address::
This parameter specifies the connection endpoint information for the data
server's data connection.

``[-data-con-port <integer>]`` - Data Server Connect Port::
This parameter specifies the TCP/IP port that the data server will use
when establishing a data connection.

``[-data-bytes-processed <integer>]`` - Data Bytes Processed::
This parameter represents the cumulative number of data stream bytes
transferred between the backup or recovery method and the data connection
during the current data operation.

``[-mover-state <component state>]`` - Mover State::
This parameter identifies the current state of the NDMP tape server's
mover state machine.

``[-mover-mode <mover mode>]`` - Mover Mode::
This parameter identifies the direction of the mover data transfer.

``[-mover-pause-reason <pause reason>]`` - Mover Pause Reason::
This parameter identifies the event that caused the mover state machine to
enter the PAUSED state.

``[-mover-halt-reason <halt reason>]`` - Mover Halt Reason::
This parameter identifies the event that caused the mover state machine to
enter the HALTED state.

``[-mover-record-size <integer>]`` - Mover Record Size::
This parameter represents the current mover record size in bytes.

``[-mover-record-num <integer>]`` - Mover Record Number::
This parameter represents the last tape record processed by the mover.

``[-mover-bytes-moved <integer>]`` - Mover Bytes Moved::
```

```
This parameter represents the cumulative number of data stream bytes
written to the data connection or the number of data stream bytes read
from the data connection and written to the tape subsystem, depending on
the mode of mover operation.

``[-mover-seek-position <integer>]`` - Mover Seek Position::
This parameter represents the data stream offset of the first byte the DMA
requested the mover to transfer to the data connection during a mover read
operation.

``[-mover-bytes-left-to-read <integer>]`` - Mover Bytes Left to Read::
This parameter represents the number of data bytes remaining to be
transferred to the data connection to satisfy the current NDMP_MOVER_READ
request.

``[-mover-window-offset <integer>]`` - Mover Window Offset::
This parameter represents the absolute offset of the first byte of the
mover window within the overall data stream.

``[-mover-window-length <integer>]`` - Mover Window Length::
This parameter represents the length of the current mover window in bytes.

``[-mover-position <integer>]`` - Mover Position::
This parameter can be used to list only those sessions, whose mover
position matches a specific value. Mover-position should be an integer.

``[-mover-setrecordsize-flag {true|false}]`` - Mover SetRecordSize Flag::
This parameter is used by the DMA to establish the record size used for
mover-initiated tape read and write operations.

``[-mover-setwindow-flag {true|false}]`` - Mover SetWindow Flag::
This parameter represents whether a mover window has been set or not. A
mover window represents the portion of the overall backup stream that is
accessible to the mover without intervening DMA tape manipulation.

``[-mover-con-addr-type <address type>]`` - Mover Connect Type::
This parameter specifies the type of data connection established by the
mover. The data connection can be established locally within a given
system or between remote networked systems.

``[-mover-con-addr <text>]`` - Mover Connect Address::
This parameter specifies the endpoint address or addresses that the mover
will use when establishing a data connection.

``[-mover-con-port <integer>]`` - Mover Connect Port::
This parameter specifies the TCP/IP port that the mover will use when
establishing a data connection.
```

```
``[-eff-host <host type>]`` - Effective Host::
This parameter indicates the host context in which the NDMP session runs.
The valid values are: PRIMARY or PARTNER.

``[-client-addr <text>]`` - NDMP Client Address::
This parameter specifies the client's IP address.

``[-client-port <integer>]`` - NDMP Client Port::
This parameter specifies the client's port number.

``[-spt-device-id <text>]`` - SCSI Device ID::
This parameter specifies the SCSI device ID.

``[-spt-ha <integer>]`` - SCSI Host Adapter::
This parameter specifies the SCSI host adapter.

``[-spt-scsi-id <integer>]`` - SCSI Target ID::
This parameter specifies the SCSI target.

``[-spt-scsi-lun <integer>]`` - SCSI LUN ID::
This parameter specifies the SCSI LUN ID.

``[-tape-device <text>]`` - Tape Device::
This parameter specifies the name to identify the tape device.

``[-tape-mode <mover mode>]`` - Tape Mode::
This parameter specifies the mode in which tapes are opened.

``[-is-secure-control-connection {true|false}]`` - Is Secure Control
Connection::
This parameter specifies whether the control connection is secure or not.

``[-data-backup-mode <text>]`` - Data Backup Mode::
This parameter specifies whether the mode of data backup is Dump or
SMTape.

``[-data-path <text>]`` - Data Path::
This parameter specifies the path of data being backed up.

``[-source-addr <text>]`` - NDMP Source Address::
This parameter specifies the control connection IP address of the NDMP
session.

== Examples

The following example displays diagnostic information about all the
```

```
sessions in the cluster:
```

cluster1::> system services ndmp probe
Node: cluster1-01
Session identifier: 4952
NDMP version: 4
Session authorized: true
Data state: IDLE
Data operation: NOACTION
Data server halt reason: NA
Data server connect type: LOCAL
....
...

```
                    Node: cluster1-02
      Session identifier: 5289
            NDMP version: 4
      Session authorized: true
              Data state: IDLE
          Data operation: NOACTION
 Data server halt reason: NA
Data server connect type: LOCAL
 ....
  ...
```

```
The following example displays diagnostic information of sessions running
on the node cluster1-01 only:
```

cluster1::> system services ndmp probe -node cluster1-01
Node: cluster1-01
Session identifier: 4952
NDMP version: 4
Session authorized: true
Data state: IDLE
Data operation: NOACTION
Data server halt reason: NA
Data server connect type: LOCAL
....
...

```
[[ID14fb024851c99a72f3643731c4cde0b2]]
= system services ndmp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
(DEPRECATED)-Display NDMP service configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "link:vserver-services-ndmp-show.html[vserver services ndmp
show] " command.

The `system services ndmp show` command displays the following information
about the NDMP configuration across all the nodes in the cluster:

* Node name
* Whether NDMP is enabled on the node
* Whether sending the NDMP password in clear text is enabled on the node
* NDMP user ID

A combination of parameters can be optionally supplied to filter the
results based on specific criteria.

== Parameters

{ ``[-fields <fieldname>,...]``::
If this parameter is specified, the command displays only the fields that
you specify.

| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed information
about all entries.

``[-node {<nodename>|local}]`` - Node::
Selects information about the specified node.

``[-enable {true|false}]`` - NDMP Service Enabled::
Selects information about the nodes where NDMP is enabled/disabled.

``[-clear-text {true|false}]`` - Allow Clear Text Password::
Selects information about the nodes whose clear-text setting matches the
specified value.

``[-user-id <text>]`` - NDMP User ID::
```

```
Selects information about the nodes that have the specified NDMP user ID.

== Examples

The following example displays information about the NDMP configuration of
all nodes in the cluster:
```

cluster1::> system services ndmp show
Node Enabled Clear Text User ID
---------- --------- ---------- -------
node0 true true ndmp
node1 true true ndmp
node2 true true ndmp
node3 true true ndmp
4 entries were displayed.

```
== Related Links

* xref:{relative_path}vserver-services-ndmp-show.html[vserver services
ndmp show]




[[ID630afe0508df7b7d022e5b267dece363]]
= system services ndmp status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of NDMP sessions

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system services ndmp status` command lists all the NDMP sessions in
the cluster. By default it lists the following details about the active
sessions:

* Node
* Session ID

A combination of parameters can be optionally supplied so as to list only
```

```
those sessions which match specific conditions. A short description of
each of the parameter is provided in the parameters section.

== Parameters

{ ``[-fields <fieldname>,...]``::
This optional parameter specifies which all additional fields to display.
Any combination of the following fields are valid:
+
  * ndmp-version
  * session-authorized
  * data-state
  * data-operation
  * data-halt-reason
  * data-con-addr-type
  * data-con-addr
  * data-con-port
  * data-bytes-processed
  * mover-state
  * mover-mode
  * mover-pause-reason
  * mover-halt-reason
  * mover-record-size
  * mover-record-num
  * mover-bytes-moved
  * mover-seek-position
  * mover-bytes-left-to-read
  * mover-window-offset
  * mover-window-length
  * mover-position
  * mover-setrecordsize-flag
  * mover-setwindow-flag
  * mover-con-addr-type
  * mover-con-addr
  * mover-con-port
  * eff-host
  * client-addr
  * client-port
  * spt-device-id
  * spt-ha
  * spt-scsi-id
  * spt-scsi-lun
  * tape-device
  * tape-modes
  * is-secure-control-connection
  * data-backup-mode
```

```
  * data-path
  * source-addr



| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed information
about all the active sessions.

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays information about the
sessions running on the specified node only. Node should be a valid node
name.

``[-session-id <integer>]`` - Session Identifier::
If this parameter is specified, the command displays information about
specific NDMP session. A session-id is a number used to identify a
particular NDMP session.

``[-ndmp-version <integer>]`` - NDMP Version::
This parameter refers to the NDMP protocol version being used in the
session.

``[-session-authorized {true|false}]`` - Session Authorized::
This field indicates whether an NDMP session is authenticated or not.

``[-data-state <component state>]`` - Data State::
This field identifies the current state of the data server's state
machine.

``[-data-operation <data operation>]`` - Data Operation::
This field identifies the data server's current operation.

``[-data-halt-reason <halt reason>]`` - Data Server Halt Reason::
This field identifies the event that caused the data server state machine
to enter the HALTED state.

``[-data-con-addr-type <address type>]`` - Data Server Connect Type::
This field specifies the type of data connection established by the data
server. The data connection can be established locally within a given
system or between remote networked systems.

``[-data-con-addr <text>]`` - Data Server Connect Address::
This specifies the connection endpoint information for the data server's
data connection.

``[-data-con-port <integer>]`` - Data Server Connect Port::
```

```
This specifies the TCP/IP port that the data server will use when
establishing a data connection.

``[-data-bytes-processed <integer>]`` - Data Bytes Processed::
This field represents the cumulative number of data stream bytes
transferred between the backup or recovery method and the data connection
during the current data operation.

``[-mover-state <component state>]`` - Mover State::
This parameter identifies the current state of the NDMP tape server's
mover state machine.

``[-mover-mode <mover mode>]`` - Mover Mode::
This parameter identifies the direction of the mover data transfer.

``[-mover-pause-reason <pause reason>]`` - Mover Pause Reason::
This parameter identifies the event that caused the mover state machine to
enter the PAUSED state.

``[-mover-halt-reason <halt reason>]`` - Mover Halt Reason::
This integer field identifies the event that caused the mover state
machine to enter the HALTED state.

``[-mover-record-size <integer>]`` - Mover Record Size::
This field represents the current mover record size in bytes.

``[-mover-record-num <integer>]`` - Mover Record Number::
This field represents the last tape record processed by the mover.

``[-mover-bytes-moved <integer>]`` - Mover Bytes Moved::
This field represents the cumulative number of data stream bytes written
to the data connection or the number of data stream bytes read from the
data connection and written to the tape subsystem, depending on the mode
of mover operation.

``[-mover-seek-position <integer>]`` - Mover Seek Position::
This field represents the data stream offset of the first byte the DMA
requested the mover to transfer to the data connection during a mover read
operation.

``[-mover-bytes-left-to-read <integer>]`` - Mover Bytes Left to Read::
This field represents the number of data bytes remaining to be transferred
to the data connection to satisfy the current NDMP_MOVER_READ request.

``[-mover-window-offset <integer>]`` - Mover Window Offset::
This field represents the absolute offset of the first byte of the mover
window within the overall data stream.
```

```
``[-mover-window-length <integer>]`` - Mover Window Length::
This field represents the length of the current mover window in bytes.

``[-mover-position <integer>]`` - Mover Position::
This parameter can be used to list only those sessions, whose mover
position matches a specific value. Mover-position should be an integer.

``[-mover-setrecordsize-flag {true|false}]`` - Mover SetRecordSize Flag::
This field is used by the DMA to establish the record size used for mover-
initiated tape read and write operations.

``[-mover-setwindow-flag {true|false}]`` - Mover SetWindow Flag::
This flag represents whether a mover window has been set or not. A mover
window represents the portion of the overall backup stream that is
accessible to the mover without intervening DMA tape manipulation.

``[-mover-con-addr-type <address type>]`` - Mover Connect Type::
This field specifies the type of data connection established by the mover.
The data connection can be established locally within a given system or
between remote networked systems.

``[-mover-con-addr <text>]`` - Mover Connect Address::
This specifies the endpoint address or addresses that the mover will use
when establishing a data connection.

``[-mover-con-port <integer>]`` - Mover Connect Port::
This specifies the TCP/IP port that the mover will use when establishing a
data connection.

``[-eff-host <host type>]`` - Effective Host::
This field indicates the host context in which the NDMP session runs. The
valid values are: PRIMARY or PARTNER.

``[-client-addr <text>]`` - NDMP Client Address::
This parameter specifies the client's IP address.

``[-client-port <integer>]`` - NDMP Client Port::
This parameter specifies the client's port number.

``[-spt-device-id <text>]`` - SCSI Device ID::
This parameter specifies the SCSI device ID.

``[-spt-ha <integer>]`` - SCSI Host Adapter::
This parameter specifies the SCSI host adapter.

``[-spt-scsi-id <integer>]`` - SCSI Target ID::
```

```
This parameter specifies the SCSI target.

``[-spt-scsi-lun <integer>]`` - SCSI LUN ID::
This parameter specifies the SCSI LUN ID.

``[-tape-device <text>]`` - Tape Device::
This parameter specifies the name to identify the tape device.

``[-tape-mode <mover mode>]`` - Tape Mode::
This parameter specifies the mode in which tapes are opened.

``[-is-secure-control-connection {true|false}]`` - Is Secure Control
Connection::
This parameter specifies whether the control connection is secure or not.

``[-data-backup-mode <text>]`` - Data Backup Mode::
This parameter specifies whether the mode of data backup is Dump or
SMTape.

``[-data-path <text>]`` - Data Path::
This parameter specifies the path of data being backed up.

``[-source-addr <text>]`` - NDMP Source Address::
This parameter specifies the control connection IP address of the NDMP
session.

== Examples

The following example displays all the NDMP sessions on the cluster:
```

cluster1::> system services ndmp status
Session
Node Id
--- ---
node-01 17479
node-01 19769
node-02 21118
3 entries were displayed.

```
The following example shows how to display only the sessions running on
node-01:
```

cluster1::> system services ndmp status -node node-01
Session
Node Id
--- ---
node-01 17479

node-01 19769
2 entries were displayed.

```
[[IDaccad7ecb077cb66c32218dbc2baf126]]
= system services ndmp log start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Start logging for the specified NDMP session

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "link:vserver-services-ndmp-log-start.html[vserver services
ndmp log start] " command.

This command is used to start logging on an active NDMP session on a node.
You can start logging two different kinds of sessions. The NDMP
``_server_`` session manages all NDMP tasks on the node. If you want to
log information regarding the NDMP server, use ``_server_`` with the `-
session-id` parameter to enable logging. If you want to log information
about a particular NDMP session, for example a restore operation, then
determine the session ID for the session using the "system services ndmp
status" command and use that ID with the `-session-id` parameter to enable
logging.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node.

``-session-id {<integer>|server}`` - Session Identifier [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the NDMP session-id on which logging needs to be
started. The session-id is associated with a unique NDMP session. Specify
``_server_`` to start logging on the NDMP server session.
```

```
``-filter <text>`` - Level Filter [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the filter for a particular session ID. This
parameter controls the NDMP modules for which logging is to be enabled.
This parameter can take five values. They are as follow : ``_all_`` ,
``_none_`` , ``_normal_`` , ``_backend_`` or ``_"filter-expression"_`` .
The default value for this is ``_none_`` .
+
  * ``_all_`` turns on logging for all modules.
  * ``_none_`` disables logging for all modules.
  * ``_normal_`` is a short cut parameter that enables logging for all
modules except ``_verbose_`` and ``_io_loop_`` . The equivalent filter
string is ``_all-verbose-io_loop_``
  * ``_backend_`` is a short cut parameter that enables logging for all
modules except ``_verbose_`` , ``_io_loop_`` , ``_ndmps_`` and ``_ndmpd_``
. The equivalent filter string is ``_all-verbose-io_loop-ndmps-ndmpp_``
  * ``_(filter-expression)_`` is a combination of one or more modules for
which logs needs to be enabled. Multiple module names can be combined
using following operators :
+
    * ``_-_`` to remove the given module from the list of specified
modules in the filter string. For example the filter ``_all-ndmpp_`` will
enable logging for all modules but not ``_ndmpp_`` .
    * ``_^_`` to add the given module or modules to the list of modules
specified in the filter string. For example the filter
``_ndmpp^mover^data_`` will enable logging for ``_ndmpp_`` , ``_mover_``
and ``_data_`` .




+
The possible module names and a brief description is given below:
```

```
+-----------------------------------------------------+
| Modules     |     Description                       |
+-----------------------------------------------------+
| verbose     | verbose message                       |
| io          | I/O process loop                      |
| io_loop     | I/O process loop verbose messages     |
| ndmps       | NDMP service                          |
| ndmpp       | NDMP Protocol                         |
| rpc         | General RPC service                   |
| fdc_rpc     | RPC to FC driver service              |
| auth        | Authentication                        |
| mover       | NDMP MOVER (tape I/O)                 |
| data        | NDMP DATA (backup/restore)            |
| scsi        | NDMP SCSI (robot/tape ops)            |
| bkup_rpc    | RPC to Backup service client          |
| bkup_rpc_s  | RPC to Backup service server          |
| cleaner     | Backup/Mover session cleaner          |
| conf        | Debug configure/reconfigure           |
| dblade      | Dblade specific messages              |
| timer       | NDMP server timeout messages          |
| vldb        | VLDB service                          |
| smf         | SMF Gateway messages                  |
| vol         | VOL OPS service                       |
| sv          | SnapVault NDMP extension              |
| common      | NDMP common state                     |
| ext         | NDMP extensions messages              |
| sm          | SnapMirror NDMP extension             |
| ndmprpc     | NDMP Mhost RPC server                 |
+-----------------------------------------------------+
```

```
== Examples

The following example shows how to start logging on a specific NDMP
session 33522, running on the node cluster1-01 with filter normal.
```

cluster1::*> system services ndmp log start -node cluster1-01 -session-id 33522 -filter normal

```
The following example shows how to start logging on the NDMP server
session, on the node cluster1-01 with filter all.
```

cluster1::*> system services ndmp log start -session-id server -filter all -node cluster1-01

```
== Related Links

* xref:{relative_path}vserver-services-ndmp-log-start.html[vserver
services ndmp log start]




[[IDf19460ccb4cb18533c1a9dd30ceea942]]
= system services ndmp log stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Stop logging for the specified NDMP session

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "link:vserver-services-ndmp-log-stop.html[vserver services
ndmp log stop] " command.

This command is used to stop logging on an active NDMP session on a node.
The NDMP ``_server_`` session manages all NDMP tasks on the node. If you
want to stop logging information regarding the NDMP server, use
``_server_`` with the `-session-id` parameter to disable logging. If you
want to stop logging information about a particular NDMP session, for
example a restore operation, then determine the session ID for the session
using the "system services ndmp status" command and use that ID with the
`-session-id` parameter to disable logging.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the node.

``-session-id {<integer>|server}`` - Session Identifier [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the NDMP session-id on which logging needs to be
```

```
stopped. The session-id is associated with a unique NDMP session. Specify
``_server_`` to stop logging on the NDMP server session.

== Examples

The following example shows how to stop logging on a specific NDMP session
35512, running on node cluster1-01.
```

cluster1::*> system services ndmp log stop -session-id 35512 -node cluster1-01

```
The following example shows how to stop logging on the NDMP server
session, running on node cluster1-01.
```

cluster1::*> system services ndmp log stop -session-id server -node cluster1-01

```
== Related Links

* xref:{relative_path}vserver-services-ndmp-log-stop.html[vserver services
ndmp log stop]




[[IDf65eb8b142d8fee36329a11fb34d487c]]
= system services ndmp node-scope-mode off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Disable NDMP node-scope-mode

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "`vserver services ndmp` " command.

This command puts NDMP server in Vserver-aware mode. The Vserver-aware
commands are available under `vserver services ndmp` .

== Examples

The following example shows how to disable the node-scope-mode of NDMP
server.
```

cluster1::> system services ndmp node-scope-mode off
NDMP node-scope-mode is disabled.

```
[[ID7af69be1ba03cd270027ca7bbbab0999]]
= system services ndmp node-scope-mode on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Enable NDMP node-scope-mode

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "`vserver services ndmp` " command.

This command puts the NDMP server in the node-scope-mode. In the node-
scope-mode, NDMP server has the following behavior:

* All NDMP operations are restricted to resources on the node
* Vserver-aware NDMP commands are disabled
* NDMP authentication falls back to ONTAP 8.1 NDMP authentication scheme



== Examples

The following example enables node-scope-mode of operation :
```

cluster1::> system services ndmp node-scope-mode on
NDMP node-scope-mode is enabled.

```
[[IDbc916086acfe49c29761d8b3dda10fa3]]
= system services ndmp node-scope-mode status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Status of NDMP node-scope-mode

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description


NOTE: This node-scoped NDMP command is deprecated. Node-scoped NDMP
functionality may be removed in a future release of ONTAP. Use the
Vserver-aware "`vserver services ndmp` " command.

This command displays whether the NDMP server is operating in node-scope-
mode or not.

* NDMP node-scope-mode is disabled - NDMP server is Vserver-aware
* NDMP node-scope-mode is enabled - NDMP server is node scoped



== Parameters

== Examples

The following example shows how to check the status of NDMP server in a
cluster
```

cluster1::> system services ndmp node-scope-mode status
NDMP node-scope-mode is disabled.

```
[[ID714b485da18cd783438e8ad6049dc564]]
= system services ndmp service modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify NDMP service configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system services ndmp service modify` command allows you to modify the
NDMP service configurations for a node in the cluster. The following
configuration can be modified:

* NDMP Common Sessions



== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the node whose NDMP configuration is to be modified.

``[-common-sessions <integer>]`` - NDMP Common Sessions [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the number of extra common NDMP sessions
supported, in addition to the number of backup and restore sessions
supported for a platform. The default value is 4 for all platforms. The
number of backup and restore sessions are platform dependent.
+
NOTE: Increasing this parameter can make the storage system unresponsive.
+


== Examples

The following example modifies the NDMP configuration on a node named
node1. The configuration sets the NDMP Common Sessions to 16:
```

cluster1::> system services ndmp modify -node node1
-common-sessions 16

```
[[IDacfc8955ed851004cb99e9be6db9438b]]
= system services ndmp service show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NDMP service configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system services ndmp service show` command displays the following
information about the NDMP service configuration across all the nodes in
the cluster:

* Node name
* NDMP Common Sessions

A combination of parameters can be optionally supplied to filter the
results based on specific criteria.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
Selects information about the specified node.

``[-common-sessions <integer>]`` - NDMP Common Sessions [.advaced-
priv]#(privilege: advanced)#::
Selects information about the nodes that have the specified number of NDMP
```

```
common sessions.

== Examples

The following example displays information about the NDMP configuration of
all nodes in the cluster:
```

cluster1::> system services ndmp service show
Node Common Sessions
---------- ----------
node0 16
node1 16
node2 16
node3 16
4 entries were displayed.

```
[[ID4b578164bf7acaf81be238c5d14686c4]]
= system services ndmp service start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start the NDMP service

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system services ndmp service start` command starts the NDMP service
daemon for a node. This is different from the xref:{relative_path}system-
services-ndmp-on.html[system services ndmp on] command. The
xref:{relative_path}system-services-ndmp-on.html[system services ndmp on]
command enables the daemon to accept NDMP requests. The NDMP service
daemon starts automatically on a node when it boots up. Use this command
to start the NDMP service daemon that has been stopped by the
xref:{relative_path}system-services-ndmp-service-stop.html[system services
ndmp service stop] command.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
The node on which the NDMP service needs to be started.

== Examples
```

cluster1::*> system services ndmp service start -node node0

```
Starts the NDMP service on node0.

== Related Links

* xref:{relative_path}system-services-ndmp-on.html[system services ndmp
on]
* xref:{relative_path}system-services-ndmp-service-stop.html[system
services ndmp service stop]




[[ID247c97e25278013e32e210e406337869]]
= system services ndmp service stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop the NDMP service

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system services ndmp service stop` command stops the NDMP service
daemon on a node. This is a disruptive command and should not be used in
normal scenarios. Processing of active sessions continues but the ability
to view or kill sessions is lost. This is different from the
xref:{relative_path}system-services-ndmp-off.html[system services ndmp
off] command. The xref:{relative_path}system-services-ndmp-off.html[system
services ndmp off] command disables new NDMP connections on the node but
does not stop the NDMP service daemon.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
The node on which the NDMP service needs to be stopped.

== Examples
```

cluster1::*> system services ndmp service stop -node node0

```
Stops the NDMP service on node0.

== Related Links

* xref:{relative_path}system-services-ndmp-off.html[system services ndmp
off]




[[IDe1fad0cc5b22be9a9738ef7e8f3d2d2d]]
= system services ndmp service terminate
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Terminate all NDMP sessions

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system services ndmp service terminate` command terminates all active
sessions on the node. This command forcefully terminates all NDMP sessions
without an opportunity for a graceful shutdown. Use
xref:{relative_path}system-services-ndmp-kill-all.html[system services
ndmp kill-all] for a clean termination of all active sessions on a node.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
The node on which the NDMP sessions need to be terminated

== Examples
```

cluster1::*> system services ndmp service terminate -node node0

```
Terminates all active NDMP sessions on node0.

== Related Links
```

```
* xref:{relative_path}system-services-ndmp-kill-all.html[system services
ndmp kill-all]




[[IDa898e909e338481235bfb4b3cdb4c8b6]]
= system services web modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the cluster-level configuration of web protocols

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command modifies the overall availability of web services in the
cluster, including the core protocol configurations for those services. In
a pre-root or unclustered scenario, its scope applies to the local node.

== Parameters

``[-external {true|false}]`` - External Web Services::
Defines whether remote clients can access HTTP or HTTPS service content.
Along with the `system services firewall` configuration, this parameter
controls the visibility for client connections. The default value for this
parameter after installation is 'true', which exports web protocols for
remote access. If no value is provided during modification, its behavior
does not change.

``[-per-address-limit <integer>]`` - Per Address Limit [.advaced-
priv]#(privilege: advanced)#::
Limits the number of connections that can be processed concurrently from
the same remote address. If more connections are accepted, those in excess
of the limit are delayed and processed after the number of connections
being processed drops below the limit. The default value is 96.

``[-http-enabled {true|false}]`` - HTTP Enabled [.advaced-
priv]#(privilege: advanced)#::
Defines whether HTTP is enabled. The default value for this parameter is
``_false_`` .
```

```
``[-csrf-protection-enabled {true|false}]`` - CSRF Protection Enabled
[.advaced-priv]#(privilege: advanced)#::
Defines whether CSRF protection is enabled. The default value is
``_true_`` .

``[-csrf-token-concurrent-limit <integer>]`` - Maximum Number of
Concurrent CSRF Tokens [.advaced-priv]#(privilege: advanced)#::
Defines how many concurrent CSRF tokens can exist at any given time. The
default value is ``_500_`` .

``[-csrf-token-idle-timeout <integer>]`` - CSRF Token Idle Timeout
(Seconds) [.advaced-priv]#(privilege: advanced)#::
Defines how long (in seconds) an unused CSRF token will exist until it
expires. The default value is ``_900_`` seconds (15 minutes).

``[-csrf-token-absolute-timeout <integer>]`` - CSRF Token Absolute Timeout
(Seconds) [.advaced-priv]#(privilege: advanced)#::
Defines how long (in seconds) a CSRF token can exist regardless of usage.
The default value is ``_0/undefined_`` , which means that it will never
time out.

``[-lif-service-policy-enforced {true|false}]`` - Enforce Network
Interface Service-Policy [.advaced-priv]#(privilege: advanced)#::
Defines whether to enforce the network interface service-policy for web
services.

== Examples

The following command changes the maximum size of the wait queue:
```

cluster1::> system services web modify -wait-queue-capacity 256

```
[[IDcf1af1c6116f8e736b16485026ceb4e5]]
= system services web show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the cluster-level configuration of web protocols

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
This command displays the overall availability of web services in the
cluster, including the core protocol configurations for those services. In
a pre-root or unclustered scenario, its output applies to the local node.
The following information explains the `External Web Services` and
`Status` attributes, two features of web services' availability.

The `External Web Services` field indicates whether remote clients are
allowed to access the HTTP or HTTPS service content. Along with the
`system services firewall` configuration, the `External Web Services`
field indicates the visibility for client connections.

The `Status` field describes the aggregated operational state of cluster-
level web services as retrieved from the `system services web node`
command. The `Status` field does not reflect whether the protocols are
externally visible, but whether the server processes are running
correctly. For detailed information about individual servers, use the
xref:{relative_path}system-services-web-node-show.html[system services web
node show] command. The following are the possible values for the `Status`
in node configuration or availability:

* online, all web services are consistently configured and working
correctly.
* partial, one or more nodes' web services are unavailable due to an error
condition.
* mixed, the nodes in the cluster do not share the same web services
configuration. This situation might occur if individual nodes were
reconfigured with the `system services web node` command.
* offline, all of the nodes' web services are unavailable due to an error
condition.
* unclustered, the current node is not part of an active cluster.



The `HTTP Enabled` field indicates whether HTTP is enabled.

The `per-address-limit` field is the limit of the number of connections
that can be processed concurrently from the same remote address. If more
connections are accepted, those in excess of the limit are delayed and
processed after the number of connections being processed drops below the
limit.

== Examples

The following example displays the availability of web services for the
cluster.
```

cluster1::> system services web show
External Web Services: true
Status: online
HTTP Protocol Port: 80
HTTPS Protocol Port: 443
HTTP Enabled: true

```
== Related Links

* xref:{relative_path}system-services-web-node-show.html[system services
web node show]




[[IDfedb49ae4e69ca76cec9bbb3938ab267]]
= system services web node show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the status of the web servers at the node level

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays operational configuration for the web server
processes on the nodes in the cluster. This output is aggregated to
produce the content for the xref:{relative_path}system-services-web-
show.html[system services web show] command.

== 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>|local}]`` - Node::
```

```
Selects the nodes that match this parameter value. Identifies the node
where the web server process is being executed.

``[-external {true|false}]`` - External Web Services::
Selects the nodes that match this parameter value. Defines whether remote
clients can access the HTTP or HTTPS service content. Along with the
`system services firewall` command configuration, this parameter controls
the visibility for client connections. The default value for this
parameter after installation is ``_true_`` , which exports web protocols
for remote access.

``[-http-port <integer>]`` - HTTP Port::
Selects the nodes that match this parameter value. Defines the HTTP port
for the node-level web services.

``[-https-port <integer>]`` - HTTPS Port::
Selects the nodes that match this parameter value. Defines the encrypted
HTTP (HTTPS) port for the node-level web services.

``[-http-enabled {true|false}]`` - HTTP Enabled::
Selects the nodes that match this parameter value. Defines whether HTTP is
enabled.

``[-per-address-limit <integer>]`` - Per Address Limit [.advaced-
priv]#(privilege: advanced)#::
Selects the nodes that match this parameter value. Limits the number of
connections that can be processed concurrently from the same remote
address. If more connections are accepted, those in excess of the limit
are delayed and processed after the number of connections being processed
drops below the limit.

``[-status {offline|partial|mixed|online|unclustered}]`` - Protocol
Status::
Selects the nodes that match this parameter value. Describes the
operational state of node-level web services. This parameter does not
reflect whether protocols are externally visible, but whether the server
processes are running correctly. The following are the possible values
that describe the service availability:
+
  * online, indicates that web services are working correctly.
  * offline, indicates that web services are unavailable due to an error
condition.
  * unclustered, indicates that the current node is not part of an active
cluster.
```

```
``[-total-hits <integer>]`` - Total HTTP Requests::
Selects the nodes that match this parameter value. Indicates the total
number of requests serviced by the web server.

``[-total-bytes <integer>]`` - Total Bytes Served::
Selects the nodes that match this parameter value. Indicates the total
number of bytes returned by the web server.

``[-lif-service-policy-enforced {true|false}]`` - Enforce Network
Interface Service-Policy::
Selects the nodes that match this parameter value. Defines whether network
interface service-policy is enforced.

== Examples

The following example displays the status of web servers for nodes in the
cluster.
```

cluster1::system services web node> show
HTTP HTTP HTTPS Total Total
Node External enabled Port Port Status HTTP Requests Bytes Served
--- --- --- --- --- --- --- ---
node1 true true 80 443 online 5 1362
node2 true true 80 443 online 5 1362
2 entries were displayed.

```
== Related Links

* xref:{relative_path}system-services-web-show.html[system services web
show]




[[IDd23f59dc024a280d0a8d5016060315f4]]
= system services web ontapi modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Unsuspend ONTAPI after auto suspend

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command allows users to un-suspend ONTAPI after it has been
automatically suspended

== Parameters

``[-suspended {true|false}]`` - ONTAPI suspended status::
Defines whether ONTAPI is currently suspended or not.

== Examples

1. Example status of ONTAPI being suspended:
```

cluster1::> system services web ontapi show
ONTAPI Suspended
---
true

```
2. The following example un-suspend ONTAPI:
```

cluster1::> system services web ontapi modify -suspended false

```
3. Attempting to manually set the suspend status to true will result in an
error:
```

cluster1::> system services web ontapi modify -suspended true

Error: command failed: ONTAPI cannot be manually suspended. Use the "vserver services web" CLI command to disable the ONTAPI service.

```
[[ID7e6fe94384564f0a9ddbc453f2a08b67]]
= system services web ontapi show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the current ONTAPI status (suspended/unsuspended)

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays if ONTAPI is currently suspended.

== Examples

The following example displays the status of ONTAPI
```

cluster1::*> system services web ontapi show
ONTAPI Suspended
---
true

```
:leveloffset: -1


= system smtape commands

:leveloffset: +1


[[ID92e1e17aea604cdc776e32e10977777f]]
= system smtape abort
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Abort an active SMTape session

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command aborts the backup or restore operations based on the session
identifier. You can perform SMTape operations using the
xref:{relative_path}system-smtape-backup.html[system smtape backup] or
xref:{relative_path}system-smtape-restore.html[system smtape restore]
commands. A unique session identifier is assigned for each new SMTape
operation. This command aborts sessions that are in active and waiting
states.

== Parameters

``-session <Sequence Number>`` - Session Identifier::
Use this parameter to specify the session identifier for a backup or
restore session.

== Examples

Abort the SMTape session with the session identifier ``_20_``
```

cluster1::> system smtape abort -session 20
Abort posted to session 20.

```
== Related Links
```

```
* xref:{relative_path}system-smtape-backup.html[system smtape backup]
* xref:{relative_path}system-smtape-restore.html[system smtape restore]




[[ID824a9cf4371f60bb7f3c6cfcacceef5d]]
= system smtape backup
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Backup a volume to tape devices

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command performs a baseline backup of a specified volume path to a
tape device. You can use the command `system hardware tape drive show` to
view the list of tape devices in the cluster. You must specify a snapshot
name to perform an SMTape backup operation. The snapshot name specified is
used as the base snapshot. A new unique session ID is assigned for this
SMTape operation and the status of the session can be monitored using the
command `system smtape status` . This session ID can be subsequently used
to perform other operations such as to find the SMTape status, abort an
SMTape operation, and continue an SMTape operation.

The volume and tape device must reside on the same node in the cluster.
You must retain the base snapshot created during this backup operation in
order to use this snapshot to re-establish a SnapMirror relationship upon
a restore.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Use this parameter to specify the Vserver name on which the volume is
located. You need not specify this parameter if only one cluster Vserver
exists.

``-volume <volume name>`` - Volume Name::
Use this parameter to specify the name of the volume that needs to be
backed up to tape.
```

```
``-backup-snapshot <snapshot name>`` - Snapshot Name::
Use this parameter to specify the name of the snapshot while performing an
SMTape backup operation.

``-tape </node_name/tape_device>`` - Tape Name::
Use this parameter to specify the name of the tape device which is used
for this SMTape operation. The format of the tape device name is
/`node_name` /`tape_device` , where `node_name` is the name of the cluster
node owning the tape and `tape_device` is the name of the tape device.

``[-tape-block-size <integer>]`` - Tape Record Size in KB::
Use this parameter to specify the tape record size in KB for backup and
restore operations. The tape record size is in multiples of 4KB, ranging
from 4KB to 256KB. The default tape record size is 240KB unless it is
specified.

== Examples

The following example will start the backup of a volume ``_datavol_`` in a
Vserver ``_vserver0_`` to a tape ``_rst0a_`` . Both the volume and tape
reside on the same node ``_cluster1-01_`` . The snapshot to be backed up
is ``_datavol_snapshot_`` and the tape record size has the value of 256KB.
```

cluster1::> system smtape backup -vserver vserver0 -volume datavol
-backup-snapshot datavol_snapshot -tape /cluster1-01/rst0a
-tape-block-size 256

Session 21 created successfully

```
The following example will start the backup of a volume ``_datavol_`` in a
Vserver ``_vserver0_`` to a tape ``_rst0a_`` . The volume ``_datavol_`` is
in a Vserver ``_vserver0_`` . Both the volume and tape reside on the same
node ``_cluster1-01_`` . The snapshot to be backed up is
``_datavol_snapshot_`` and the tape record size has the default value of
240KB.
```

cluster1::> system smtape backup -vserver vserver0 -volume datavol
-backup-snapshot datavol_snapshot -tape /cluster1-01/nrst0l
Session 22 created successfully

```
[[ID3977287c94615c0b5ac34100ba45bc03]]
= system smtape break
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Make a restored volume read-write

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command breaks the relationship between the tape backup of a volume
and a restored volume, changing the restored volume from read-only to
read/write.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Use this parameter to specify the Vserver name on which the volume is
located. You need not specify this parameter if only one cluster Vserver
exists.

``-volume <volume name>`` - Volume Name::
Use this parameter to specify the name of the read-only volume that needs
to be made read/writeable after a restore.

== Examples

Make the read-only volume ``_datavol_`` on Vserver ``_vserver0_``
writeable after a restore.
```

cluster1::> system smtape break -vserver vserver0 -volume datavol
[Job 84] Job succeeded: SnapMirror Break Succeeded

```
[[ID91c9f7d637a671df3e28ec23716fe5f1]]
= system smtape continue
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Continue SMTape session waiting at the end of tape

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command continues the SMTape backup and restore operations using the
specified tape device. You can use this command when an SMTape operation
has reached the end of current tape and is in the wait state to write to
or read from a new tape.

If a tape device is not specified, the original tape device will be used.

User has to make sure that the correct tape media is inserted in the
device and positioned appropriately before issuing this command.

== Parameters

``[-tape </node_name/tape_device>]`` - Tape Name::
Use this parameter to specify the name of the tape device which is used
for this SMTape operation. The format of the tape device name is
/`node_name` /`tape_device` , where `node_name` is the name of the cluster
node owning the tape and `tape_device` is the name of the tape device.

``-session <Sequence Number>`` - Session Identifier::
Use this parameter to specify the session identifier for the SMTape backup
or restore operations.

== Examples

Continues an SMTape session having session ID ``_20_`` on tape device
``_rst0a_`` on the node ``_node1_`` in the cluster.
```

cluster1::> system smtape continue -session 20 -tape /node1/rst0a
continue on session 20 succeeded

```
The following example continues session ``_40_`` on the same tape device
that was being used by the session.
```

cluster1::> system smtape continue -session 40
continue on session 40 succeeded

```
[[ID08e3e594e9d29f9ddca928e53afa3777]]
= system smtape restore
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore a volume from tape devices

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command performs restore of a backup image created using the command
xref:{relative_path}system-smtape-backup.html[system smtape backup] in the
specified tape device to a destination volume path. A new unique session
ID is assigned for this operation; the status of the session can be
monitored using the command `system smtape status` . It is required that
the volume and tape device reside in the same cluster node. The volume
must be of type DP (Data Protection) and should be placed in the
restricted mode prior to a restore.



Any existing data on the volume will get overwritten upon a restore. The
volume will remain as read-only and of type DP after the restore. You can
use the command xref:{relative_path}system-smtape-break.html[system smtape
break] to get read/write permissions on the volume. Restore can be done to
a non-root DP volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Use this parameter to specify the Vserver name on which the volume is
located. You need not specify this parameter if only one cluster Vserver
exists.
```

```
``-volume <volume name>`` - Volume Name::
Use this parameter to specify the volume name on which the tape content
will be restored.

``-tape </node_name/tape_device>`` - Tape Name::
Use this parameter to specify the name of the tape device which is used
for this SMTape operation. The format of the tape device name is
/`node_name` /`tape_device` , where `node_name` is the name of the cluster
node owning the tape and `tape_device` is the name of the tape device.

``[-tape-block-size <integer>]`` - Tape Record Size in KB::
Use this parameter to specify the tape record size in KB for backup and
restore operations. The tape record size is in multiples of 4KB, ranging
from 4KB to 256KB. The default tape record size is 240KB unless it is
specified. Use the same record size which was used during the backup. If
the tape record size is different from the tape record size that was used
at the time of backup then `system smtape restore` will fail.

== Examples

The following example will start the restore to a volume ``_datavol_``
from a tape ``_rst0a_`` . The volume ``_datavol_`` is in a Vserver
``_vserver0_`` . Both ``_vserver0_`` and ``_rst0a_`` reside on the same
node ``_cluster1-01_`` .
```

cluster1::> system smtape restore -vserver vserver0 -volume datavol
-tape /cluster1-01/rst0a -tape-block-size 256
Session 2 created successfully

```
The following example will start the restore to a volume ``_datavol_``
from a tape ``_rst0a_`` . The volume ``_datavol_`` is in a Vserver
``_vserver0_`` . Both ``_vserver0_`` and ``_rst0a_`` reside on the same
node ``_cluster1-01_`` . The default tape record size of 240KB was used
during backup.
```

cluster1::> system smtape restore -vserver vserver0 -volume datavol
-tape /cluster1-01/rst0a
Session 5 created successfully

```
== Related Links

* xref:{relative_path}system-smtape-backup.html[system smtape backup]
* xref:{relative_path}system-smtape-break.html[system smtape break]
```

```
[[ID686a067124ea929104b7511116e0be92]]
= system smtape showheader
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display SMTape header

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the image header of a tape. The tape must have a
valid backup of data. The following information about the backup is
displayed:

* Tape Number - the tape number if the backup spans multiple tape devices.
* WAFL Version - WAFL version of the storage system when the volume was
backed up on tape.
* Backup Set ID - a unique backup set ID for the baseline backup.
* Source Storage System - the source storage system where the volume
resided when the backup was performed.
* Source Volume - the source volume that was backed up to tape.
* Source Volume Capacity - the capacity of the source volume that was
backed up to tape.
* Source Volume Used Size - the used size of the source volume that was
backed up to tape.
* Source Snapshot - name of the snapshot used for the backup.
* Volume Type - type of the volume.
* Is SIS Volume - this field is true if the backed up volume was a SIS
volume.
* Backup Version - the SMTape backup version.
* Backup Sequence No - the backup sequence number.
* Backup Mode - this field describes the backup mode.
* Time of Backup - the time at which the backup was performed.
* Time of Previous Backup - the time at which the previous backup was
performed; this information is displayed only if the previous backup was
an incremental backup.
* Volume Total Inodes - number of inodes of the backed up volume.
* Volume Used Inodes - number of used inodes of the backed up volume.
* Number of Snapshots - number of snapshots present in this backup.
```

```
* Snapshot ID - is the Snapshot ID of the backup Snapshot.
* Snapshot Time - time at which the backup snapshot was created.
* Snapshot Name - name of the snapshot which was backed up to tape.



== Parameters

``-tape </node_name/tape_device>`` - Tape Name::
Use this parameter to specify the name of the tape device which is used
for this SMTape operation. The format of the tape device name is
/`node_name` /`tape_device` , where `node_name` is the name of the cluster
node owning the tape and `tape_device` is the name of the tape device.

``[-tape-block-size <integer>]`` - Tape Record Size in KB::
Use this parameter to specify the tape record size in KB for backup and
restore operations. The tape record size is in multiples of 4KB, ranging
from 4KB to 256KB. The default tape record size is 240KB unless it is
specified.

== Examples

The following example reads the image header from the tape ``_nrst0l_``
residing on the node ``_cluster1-01_`` and displays relevant tape header
information.
```

cluster1::> system smtape showheader -tape /cluster1-01/nrst0l
-tape-block-size 240
Tape record size in KB: 240
Tape Number: 1
WAFL Version: 23577
Backup Set ID: 7d0c9a15-8e20-11e1-8741-123478563412
Source Storage System: cluster1-01
Source Volume: /vs1/srcvol
Source Volume Capacity: 400.00MB
Source Volume Used Size: 0.00
Source Snapshot: mysnap
Volume Type: Flex
Is SISVolume: no
Backup Version: 1:3
Backup Sequence No: 0
Backup Mode: dw-data
Time of Backup: 4/24/2012 15:16:38
Time of Previous Backup: 0/0/0 00:00:00
Volume Total Inodes: 12789
Volume Used Inodes: 100
Number of Snapshots: 1
Snapshot ID: 1
Snapshot Time: 4/24/2012 15:16:10

Snapshot Name: mysnap

```
[[ID07deb16ec25f366e333ed8bd1b04c165]]
= system smtape status clear
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Clear SMTape sessions

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command clears SMTape sessions which are completed, failed or Unknown
state.

== Parameters

``[-session <Sequence Number>]`` - Session Identifier::
Use this parameter to clear the SMTape sessions with the specified session
identifier.

``[-node {<nodename>|local}]`` - Node Name::
Use this parameter to clear the SMtape sessions related to the specified
node.

``[-type {backup|restore}]`` - Operation Type::
Use this parameter to clear the SMTape sessions of the specified operation
type. These can be either backup or restore sessions.

``[-status {COMPLETED|FAILED|UNKNOWN}]`` - Session Status::
Use this parameter to clear the SMTape sessions which have the status as
specified in the parameter.

``[-path <text>]`` - Path Name::
Use this parameter to clear the SMTape sessions which have path as
specified in the parameter.

``[-device <text>]`` - Device Name::
Use this parameter to clear the SMTape sessions on a specific tape device.

``[-backup-snapshot <snapshot name>]`` - Snapshot Name::
```

```
Use this parameter to clear the SMTape sessions using the snapshot name as
specified in the parameter.

``[-tape-block-size <integer>]`` - Tape Block Size::
Use this parameter to clear the SMTape sessions with the tape block size
as specified in the parameter.

== Examples

The following example clears all the completed SMTape sessions in the
cluster:
```

cluster1::> system smtape status clear
5 sessions are purged.

```
The SMTape sessions on the node ``_node1_`` in the cluster are cleared.
```

cluster1::> system smtape status clear -node node1
3 sessions are purged.

```
[[ID515a9a14a5570a05da71a75d63897941]]
= system smtape status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show status of SMTape sessions

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command lists the status of all SMTape sessions in the cluster. By
default, this command lists the following information:

* Session
* Type
* Status
* Progress
* Path
* Device
* Node
```

```
== Parameters

{ ``[-fields <fieldname>,...]``::
Use this parameter to display additional fields about each session apart
from the default entries. This parameter is optional. Any combination of
the following fields is valid:
+
  * Session
  * Node
  * Type
  * Status
  * Path
  * Device
  * Progress
  * Start-time
  * End-time
  * Update-time
  * Backup-snapshot
  * Tape-block-size
  * Error



| ``[-instance  ]`` }::
Displays detailed information about the specified sessions.

``[-session <Sequence Number>]`` - Session Identifier::
Selects information about a specific SMTape session. A Session Identifier
is a number that is used to identify a particular SMTape session.

``[-node {<nodename>|local}]`` - Node Name::
Selects information about sessions related to the specified node.

``[-type {backup|restore}]`` - Operation Type::
 Selects information about SMTape sessions of the specified operation
type.

``[-status {COMPLETED|FAILED|ACTIVE|WAITING|ABORTING|UNKNOWN}]`` - Session
Status::
Selects information about SMTape sessions having the specified status in
the parameter.

``[-path <text>]`` - Path Name::
Selects information about SMTape sessions on a volume which is at the
```

```
specified path name. This is the logical path of the volume and you must
specify the path name in the following format: /`vserver_name`
/`volume_name` .

``[-device <text>]`` - Device Name::
Selects information about the SMTape sessions on the specified tape
device. You must specify the tape device name in the following format:
/`node_name` /`tape_device` .

``[-progress {<integer>[KB|MB|GB|TB|PB]}]`` - Bytes Transferred::
Selects information about SMTape sessions in which the number of data
bytes transferred in a particular session matches with the number
specified in this parameter.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Start Time::
Selects information about SMTape sessions whose starting time matches the
specified starting time.

``[-end-time <MM/DD/YYYY HH:MM:SS>]`` - End Time::
Selects information about SMTape sessions whose ending time matches the
specified ending time.

``[-backup-snapshot <snapshot name>]`` - Snapshot Name::
Selects information about SMTape sessions that use a particular snapshot
name which matches the specified snapshot name in the parameter in backup
or restore operations.

``[-tape-block-size <integer>]`` - Tape Block Size::
Selects information about SMTape sessions that use a particular tape block
size which matches the specified tape block size parameter in backup or
restore operations.

``[-error <text>]`` - Error Description::
Selects information about SMTape sessions that have a particular error
description which matches the specified error description in the
parameter.

== Examples

Displays default entries about the five SMTape sessions.
```

cluster1::> system smtape status show

Session Type Status Progress Path Device Node
------- ------- --------- -------- ------------ --------------- -----------
5 Backup COMPLETED 50MB /vsrvr1/vol1 /cls1-01/nrst0l cluster1-01
4 Restore FAILED 0B /vsrvr1/vol3 /cls1-02/nrst2l cluster1-02

3 Backup COMPLETED 50MB /vsrvr1/vol3 /cls1-01/nrst0l cluster1-01
2 Backup COMPLETED 50MB /vsrvr1/vol2 /cls1-03/nrst0m cluster1-03
1 Backup COMPLETED 50KB /vsrvr1/vol5 /cls1-01/nrst0n cluster1-01
5 entries were displayed.

```
The following example shows the output with the `-instance` argument.
```

cluster1::> system smtape status show -instance

Session Identifier: 1
Node Name: node1
Operation Type: Backup
Status: COMPLETED
Path Name: /vs1/vol1
Device Name: /node1/rst0a
Bytes Transferred: 2048
Start Time: 1/4/2012 14:26:24
End Time: 1/4/2012 14:29:45
Last updated: 1/4/2012 14:29:45
Snapshot Name: vol1.snapshot
Tape Block Size: 240
Error Description: None

```
:leveloffset: -1


= system snmp commands

:leveloffset: +1


[[IDa10e37e5b851ef92a05e9f12cf668fb5]]
= system snmp authtrap
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enables or disables SNMP authentication traps

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Use this command to either enable or disable the standard SNMP
authentication failure traps.

== Parameters

``[-authtrap <integer>]`` - Enables SNMP Authentication Trap::
Enter the value of 1 to enable SNMP authentication failure traps. By
default, SNMP authentication trap is disabled and the value is 0.

== Examples

The following example demonstrates how to set the SNMP authtrap. +
```

cluster1::> system snmp authtrap -authtrap 1
uster1::> system snmp show
contact:
private
location:
NB
authtrap:
1
init:
0
traphosts:

-
community:
- -

```
[[ID2d33c95eca90d84cc17fe149fcb4ea16]]
= system snmp contact
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays or modifies contact details

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Sets the contact name as the System.sysContact.0 MIB-II variable.

== Parameters

``[-contact <text>]`` - Contact::
Specifies the contact name. Without any value specified, this command
displays current setting of contact name.

== Examples

The following example sets the contact name for SNMP. +
```

cluster1::> system snmp contact -contact private
uster1::> system snmp show
contact:
private
location:
NB
authtrap:
1
init:
0
traphosts:
-
community:
- -

```
[[ID822a80b791058c7c72be31e4b38e8b22]]
= system snmp enable-snmpv3
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enables SNMPv3 cluster-wide

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system snmp enable-snmpv3` command enables SNMPv3 server on the
entire cluster. When this command is run, SNMP users and SNMP traphosts
that are non-compliant to FIPS will be deleted automatically, since
cluster FIPS mode is enabled. Any SNMPv1 user, SNMPv2c user or SNMPv3 user
(with none or MD5 as authentication protocol or none or DES as encryption
protocol or both) is non-compliant to FIPS. Any SNMPv1 traphost or SNMPv3
traphost (configured with an SNMPv3 user non-compliant to FIPS) is non-
compliant to FIPS.

== Examples

The following command enables SNMPv3 server on the entire cluster, within
a cluster named cluster1:
```

cluster1::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y

cluster1::*> system snmp enable-snmpv3

Warning: If you enable SNMPv3 using this command, any SNMP users and
SNMP traphosts that are non-compliant to FIPS will be
deleted automatically, since cluster FIPS mode is enabled.
Any SNMPv1 user, SNMPv2c user or SNMPv3 user (with none or
MD5 as authentication protocol or none or DES as encryption
protocol or both) is non-compliant to FIPS. Any SNMPv1
traphost or SNMPv3 traphost (configured with an SNMPv3 user
non-compliant to FIPS) is non-compliant to FIPS.
Do you want to continue? {y|n}: y

1 entry was modified.

cluster1::*>

```
[[ID441c34ac503aeed4199b24a5103c5cd4]]
= system snmp init
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enables or disables SNMP traps

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Initializes or disables sending of traps by the SNMP daemon from the
cluster.

== Parameters

``[-init <integer>]`` - Initialize Traps::
Use the value of 1 to initialize SNMP daemon to send traps or use a value
of 0 to stop sending traps from the cluster. If no value is specified,
this command displays the current setting of init. Traps are enabled by
default.

== Examples

The following command initializes SNMP daemon to send traps. +
```

cluster1::> system snmp init -init 1
uster1::> system snmp show
contact:
private
location:
NB
authtrap:
1
init:
1
traphosts:
-
community:
- -

```
[[ID54c267ed8e31b2d5571c7b2bcb86328a]]
= system snmp location
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays or modifies location information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Sets the location name as the System.sysLocation.0 MIB-II variable.

== Parameters

``[-location <text>]`` - Location::
Specifies the location details. If no value is specified, this command
displays the current setting of location.

== Examples

This command sets the location name. +
```

cluster1::> system snmp location -location NB
cluster1::> system snmp show
contact:
private
location:
NB
authtrap:
1
init:
1
traphosts:
-
community:
- -

```
[[ID6cb6bcad2a463ab64335efa21da7ad27]]
= system snmp prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change SNMP configuration to the default settings for releases earlier
than Data ONTAP 9.3.0

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system snmp prepare-to-downgrade` command prepares the SNMP subsystem
for a downgrade or a revert. More specifically, it prepares the SNMPv3
client feature for a downgrade or a revert. It deletes all storage
switches that were explicitly added for monitoring and are using SNMPv3 as
the underlying protocol. It also deletes any cluster switches that are
using SNMPv3 for monitoring. Finally, it deletes any remote switch SNMPv3
users configured in ONTAP.

== Examples

The following command prepares the SNMP subsystem for a downgrade or a
revert, within a cluster named cluster1:
```

cluster1::*> system snmp prepare-to-downgrade

```
[[ID7aa3bc949cfa8195576004e43d68ac5e]]
= system snmp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays SNMP settings

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Lists the current values of all the SNMP parameters.

== Examples

The example below shows a typical command display.
```

cluster1::> system snmp show
contact:
private
location:
NB
authtrap:
1
init:
1
traphosts:
xxx.example.com(xxx.example.com)(192.168.xxx.xxx)
community:
- -

```
[[ID11b218748d7d69b4be1564d64e7787a5]]
= system snmp community add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Adds a new community with the specified access control type

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system snmp community add` command adds communities with the
specified access control type. Only read-only communities are supported.
There is no limit for the number of communities supported.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver to which the community will be added.
If no Vserver is specified, the community is added to the admin Vserver.

``-community-name <text>`` - Community::
This parameter specifies the name of the community.

``-type <ctype>`` - access type::
This parameter specifies 'ro' for read-only community.

== Examples

The following example adds the read-only community name 'private'.
```

cluster1::> system snmp community add -type ro
-community-name private
cluster1::> system snmp community show
ro private

```
[[ID2b81246c514f5a3449757c9e04b89855]]
= system snmp community delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Deletes community with the specified access control type

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system snmp community delete` command deletes communities with the
specified access control type. Only read-only communities are supported.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver from which you want to delete the
community. If no Vserver is specified, the community is deleted from the
admin Vserver.

``-community-name <text>`` - Community::
Specify the name of the community.

``-type <ctype>`` - access type::
Specify 'ro' for a read-only community.

== Examples

The following example deletes the read-only community 'private':
```

cluster1::> system snmp community delete -type ro
-community-name private
uster1::> system snmp community show
This table is currently empty.

```
[[ID51af26f33942a9791ec7f2044e21c68b]]
= system snmp community show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays communities

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Displays the current list of SNMP communities.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Selects the Vserver to which the SNMP community belongs

``[-community-name <text>]`` - Community::
Selects the SNMP v1/v2c community string

``[-access <ctype>]`` - access::
Selects the access type of the SNMP v1/v2c community. Read-only (ro) is
the only access type supported

== Examples
```

cluster1::> system snmp community show
cluster1
ro private

```
[[IDb161de425d440d52e1b47d31296f65dc]]
= system snmp traphost add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a new traphost

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Adds the SNMP manager who receives the SNMP trap PDUs. The SNMP manager
can be a hostname or IP address. There is no limit on the number of
traphosts supported.

== Parameters

``-peer-address <Remote InetAddress>`` - Remote IP Address::
Specifies the IP address or hostname of the traphost. If the USM user is
associated, then the SNMPv3 traps are generated for this traphost using
the associated USM user's authentication and privacy credentials. If no
USM user is associated, then the SNMP v1/v2c traps are generated for this
traphost. For the SNMP v1/v2c traps, the default community string is
'public', when no community is defined. When the community strings are
defined, then the first community string is chosen for the SNMP v1/v2c
traps.

``[-usm-username <text>]`` - USM User Name::
Specifies a predefined SNMPv3 USM user. The SNMPv3 traps are generated
using this USM user's authentication and privacy credentials for the
traphost identified by the peer-address parameter.

== Examples

In the following example, the command adds a hostname 'yyy.example.com'
for the SNMPv3 traps: +
```

cluster1::> system snmp traphost add -peer-address yyy.example.com -usm-username MyUsmUser
cluster1::> system snmp traphost show
yyy.example.com(yyy.example.com)(192.168.xxx.xxx) USM User: MyUsmUser

```
In the following example, the command adds a hostname 'xxx.example.com'
for the SNMP v1/v2c traps: +
```

cluster1::> system snmp traphost add xxx.example.com
cluster1::> system snmp traphost show
yyy.example.com(yyy.example.com)(192.168.xxx.xxx) USM User: MyUsmUser
xxx.example.com(xxx.example.com)(xxx.xxx.xxx.xxx) Community: public

```
[[IDcdb16432bfbb665ee4ebe79ac1067de5]]
= system snmp traphost delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a traphost

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Deletes the SNMP manager, who receives the SNMP trap PDUs. The SNMP
manager can be a hostname or IP address. There is no limit on the number
of traphosts supported.

== Parameters

``-peer-address <Remote InetAddress>`` - Remote IP Address::
Specifies the IP address or hostname of the traphost. If the USM user is
associated, then specify the USM user to delete the traphost.

``[-usm-username <text>]`` - USM User Name::
Specifies the USM user associated with traphost.

== Examples

In the following example, the command deletes the SNMPv3 traphost
'yyy.example.com' associated with the USM user: +
```

cluster1::> system snmp traphost delete -peer-address yyy.example.com -usm-username MyUsmUser

```
In the following example, the command deletes the SNMP v1/v2c traphost
'xxx.example.com' associated with a community string: +
```

cluster1::> system snmp traphost delete -peer-address xxx.example.com

```
[[ID7fdb639153c1534519636c60818882a3]]
= system snmp traphost show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays traphosts

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Displays list of the SNMP v1/v2c and SNMP v3 managers, that receive trap
PDUs.

== Examples

In the following example, the command displays all the host names or IP
addresses that have been added until now: +
```

cluster1::> system snmp traphost show
yyy.example.com(yyy.example.com)(192.168.xxx.xxx) USM User: MyUsmUser
xxx.example.com(xxx.example.com)(xxx.xxx.xxx.xxx) Community: public

```
:leveloffset: -1


= system switch commands

:leveloffset: +1


[[IDf3a02cbaa5e1cc94ec04084c48abc619]]
= system switch ethernet configure-health-monitor
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Ethernet switch health monitor configuration file setup.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system switch ethernet configure-health-monitor` command downloads an
Ethernet switch's health monitor configuration file in the ZIP format,
which contains the XML file and a signed version file. After download,
ONTAP will check the signed file. If valid, the Ethernet switch health
monitor restarts to use the new Ethernet switch health monitor
configuration file.

== Parameters

``-node {<nodename>|local}`` - Node [.advaced-priv]#(privilege:
advanced)#::
Node on which the download and signed file validation process occurs. The
Ethernet switch health monitor process will be restarted on this node. The
configuration file will be ingested and made available to all nodes in the
cluster.

``-package-url <text>`` - Package URL [.advaced-priv]#(privilege:
advanced)#::
URL that provides the location of the package to be downloaded. Standard
URL schemes, including HTTP, HTTPS, FTP and FILE, are accepted.

== Examples

The following example downloads Ethernet switch health monitor
configuration file to node1 from a web server and enables Ethernet switch
health monitor to process it:
```

cluster1::*> system switch ethernet configure-health-monitor -node node1 -package-url http://example.com/hm_config.zip

```
[[ID21d622699ac784e0f85df1f950a7d71e]]
= system switch ethernet create
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add information about an Ethernet switch (cluster or storage).

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet create` command manually configures and
enables the health monitoring of a specified Ethernet switch of type:
cluster or storage. This command might be needed when ONTAP is unable to
automatically discover and initiate monitoring of a switch through the
Cisco Discovery Protocol (CDP) or the Link Layer Discovery Protocol
(LLDP). If ONTAP has discovered a switch that is not in
xref:{relative_path}system-switch-ethernet-show.html[system switch
ethernet show] , such as if the switch was deleted using
xref:{relative_path}system-switch-ethernet-delete.html[system switch
ethernet delete] with the `force` parameter, the `device` , `address` ,
and `type` parameters must match the discovered values. To identify which
switches are currently monitored, use the xref:{relative_path}system-
switch-ethernet-show.html[system switch ethernet show] command. If this
switch was previously monitored within the past 7 days, hourly
AutoSupport(R) log collection might automatically resume.

== Parameters

``-device <text>`` - Device Name::
Name of the switch intended for monitoring. ONTAP utilizes this device
name to recognize the SNMP agent it needs to interact with.

``-address <IP Address>`` - IP Address::
IP address of the switch's management interface. This address is used as
the destination of SNMP traffic and SSH connections for the purpose of
monitoring and log collection, respectively.

``-snmp-version {SNMPv1|SNMPv2c|SNMPv3}`` - SNMP Version::
SNMP version that ONTAP employs for sending SNMP requests for monitoring
purposes. By default, SNMPv2c is used, as established by the Reference
Configuration File (RCF) applied to the switch.

``-community-or-username <text>`` - SNMPv2c Community String or SNMPv3
Username::
Community string used for SNMPv2 authentication, or the SNMPv3 username
for SNMPv3 security. By default, the community string for SNMPv2
```

```
authentication is set to `cshm1!` , as determined by the RCF that is
applied to the switch. If SNMPv3 is used, the switch must be configured
with the SNMPv3 username.

``-model {OTHER|CN1610|NX3132V|NX3232C}`` - Model Number::
Model of the switch. Use the OTHER option when adding a switch that
requires an Ethernet switch health monitor XML configuration file, such as
the BES-53248, MSN2100-CB2FC, MSN2100-CB2RC, N9K-C92300YC, and N9K-C9336C-
FX2 switches. Be aware that the OTHER option selected during creation is
different from the OTHER displayed by the xref:{relative_path}system-
switch-ethernet-show.html[system switch ethernet show] command. In this
show command, OTHER indicates that a switch is not supported for health
monitoring.

``-type {cluster-network|storage-network}`` - Switch Network::
Switch type: `cluster-network` or `storage-network` .

``[-is-monitoring-enabled-admin {true|false}]`` - Enable Switch
Monitoring::
Monitoring status selected by the administrator, which is set to `true` by
default when not specified. During maintenance periods, a switch not
manually added might be repeatedly discovered and dropped, which could
potentially generate unnecessary alerts if monitoring is enabled by
default. In such cases, setting this parameter to `false` disables the
monitoring process.

== Examples

Example 1: Initiates Ethernet switch health monitoring for a Cisco Nexus
3132Q-V cluster switch named SwitchA.
```

cluster1::> system switch ethernet create -device SwitchA -address 1.2.3.4 -snmp-version SNMPv2c -community-or-username cshm1! -model NX3132V -type cluster-network

```
Example 2: Initiates Ethernet switch health monitoring for a NVIDIA SN2100
storage switch named SwitchB, with the SNMPv3 username snmpv3u1.
```

cluster1::> system switch ethernet create -device SwitchB -address 5.6.7.8 -snmp-version SNMPv3 -community-or-username snmpv3u1 -model OTHER -type storage-network

```
== Related Links

* xref:{relative_path}system-switch-ethernet-show.html[system switch
ethernet show]
* xref:{relative_path}system-switch-ethernet-delete.html[system switch
```

```
ethernet delete]




[[ID1e53179027de3663a272fa2d7aa518a7]]
= system switch ethernet delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete information about an Ethernet switch (cluster or storage).

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet delete` command is used to delete an Ethernet
switch. This command is intended to be used to delete a switch that has
been disconnected from the system, instead of waiting for it to be deleted
automatically. If you want to stop monitoring the switch, use
xref:{relative_path}system-switch-ethernet-modify.html[system switch
ethernet modify] to set the ``_is-monitoring-enabled-admin_`` parameter to
``_false_`` . By default, this command only allows deletion of a switch if
the ``_is-discovered_`` parameter is ``_false_`` (ONTAP isn't receiving
Cisco Discovery Protocol (CDP) or Link Layer Discovery Protocol (LLDP)
data for the switch). However, if the `force` parameter is used, the
switch will be deleted regardless of ``_is-discovered_`` . If the switch
is rediscovered, it will be re-added automatically.

== Parameters

``-device <text>`` - Device Name::
Name of the Ethernet switch that you want to delete.

``[-force <true>]`` - Force Delete [.advaced-priv]#(privilege:
advanced)#::
This parameter forces the delete operation. This allows deletion of a
switch when ``_is-discovered_`` is ``_true_`` . Note: The switch may be
re-added automatically.

== Examples

Example 1: Delete a removed switch named SwitchA.
```

```
cluster1::> system switch ethernet delete -device SwitchA
```

```
Example 2: Forcefully remove a discovered switch named SwitchB.
```

```
cluster1::*> system switch ethernet delete -device SwitchB -force
```

```
== Related Links

* xref:{relative_path}system-switch-ethernet-modify.html[system switch
ethernet modify]




[[IDe09b4787fd15e9561d4880ee8709dd76]]
= system switch ethernet modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify information about an Ethernet switch's configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet modify` command is used to modify the settings
of an Ethernet switch for health monitoring purposes. This command allows
you to adjust various parameters related to the switch, including its SNMP
version, SNMPv2c community or SNMPv3 username, and monitoring status. If
``_is-discovered_`` is ``_false_`` , the IP address and type can also be
changed.

== Parameters

``-device <text>`` - Device Name::
Name of the Ethernet switch that you want to modify. ONTAP utilizes this
device name to recognize the SNMP agent it needs to interact with.

``[-address <IP Address>]`` - IP Address::
IP address of the switch's management interface. This address is used as
the destination of SNMP traffic and SSH connections for the purpose of
```

```
monitoring and log collection, respectively.

``[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}]`` - SNMP Version::
SNMP version that ONTAP employs for sending SNMP requests for monitoring
purposes. By default, SNMPv2c is used, as established by the Reference
Configuration File (RCF) applied to the switch.

``[-community-or-username <text>]`` - SNMPv2c Community String or SNMPv3
Username::
Community string used for SNMPv2 authentication, or the SNMPv3 username
for SNMPv3 security. By default, the community string for SNMPv2
authentication is set to `cshm1!` , as determined by the RCF that is
applied to the switch. If SNMPv3 is used, the switch must be configured
with the SNMPv3 username.

``[-type {cluster-network|storage-network}]`` - Switch Network::
Switch type: `cluster-network` or `storage-network` .

``[-is-monitoring-enabled-admin {true|false}]`` - Enable Switch
Monitoring::
Monitoring status selected by the administrator, which is set to `true` by
default when not specified. During maintenance periods, unnecessary alerts
may be generated if monitoring is left enabled. In such cases, setting
this parameter to `false` disables the monitoring process.

== Examples

Example 1: Modifies the IP address for the switch named SwitchA. All other
settings of the switch are preserved.
```

cluster1::> system switch ethernet modify -device SwitchA -address 2.3.4.5

```
Example 2: Modifies the SNMP parameters for the switch named SwitchB. All
other settings are preserved.
```

cluster1::> system switch ethernet modify -device SwitchB -snmp-version SNMPv3 -community-or-username snmpv3u1

```
[[ID0884070e407d583ce3bf8d3b02b38918]]
= system switch ethernet show-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
(DEPRECATED)-Displays the list of switches

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description


NOTE: This command has been deprecated. It may be removed from a future
release of ONTAP.

The `system switch ethernet show-all` command displays configuration
details for cluster and storage Ethernet switches.

== Parameters

``[-fields <fieldname>,...]``::
Selects the fields that have the specified name.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
Selects detailed information for all the switches.

``[-device <text>]`` - Device Name [.advaced-priv]#(privilege:
advanced)#::
Selects the switches that match the specified device name.

``[-address <IP Address>]`` - IP Address [.advaced-priv]#(privilege:
advanced)#::
Selects the switches that match the specified IP address.

``[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}]`` - SNMP Version [.advaced-
priv]#(privilege: advanced)#::
Selects the switches that match the specified SNMP version.

``[-community-or-username <text>]`` - SNMPv2c Community String or SNMPv3
Username [.advaced-priv]#(privilege: advanced)#::
Selects the switches that match the specified community string or SNMPv3
username.

``[-discovered {true|false}]`` - Is Discovered [.advaced-priv]#(privilege:
advanced)#::
Selects the switches that match the specified discovery setting.

``[-type {cluster-network|storage-network}]`` - Switch Network [.advaced-
priv]#(privilege: advanced)#::
Selects the switches that match the specified switch type.
```

```
``[-sw-version <text>]`` - Software Version [.advaced-priv]#(privilege:
advanced)#::
Selects the switches that match the specified software version.

``[-is-monitoring-enabled-operational {true|false}]`` - Switch Monitoring
Status [.advaced-priv]#(privilege: advanced)#::
Selects the switches that match the specified operational monitoring
status.

``[-reason <text>]`` - Reason For Not Monitoring [.advaced-
priv]#(privilege: advanced)#::
Selects the switches that match the specified reason.

``[-version-source <text>]`` - Source Of Switch Version [.advaced-
priv]#(privilege: advanced)#::
Selects the switches that match the specified version source (for example,
SNMP, CDP, ISDP, or LLDP).

``[-rcf-version <text>]`` - Reference Config File Version [.advaced-
priv]#(privilege: advanced)#::
Selects the switches that match the specified reference configuration file
version.

``[-serial-number <text>]`` - Serial Number of the Device [.advaced-
priv]#(privilege: advanced)#::
Selects the switches that match the specified serial number.

``[-model <text>]`` - Model to display [.advaced-priv]#(privilege:
advanced)#::
Selects the switches that match the specified model number.

== Examples
```

cluster1::> system switch ethernet show-all
Switch Type Address Model
--- --- --- ---
SwitchA cluster 1.2.3.4 Nexus5010
Is Monitored: yes
Reason:
Software Version: Cisco IOS 4.1N1
Version Source: CDP

```
The example above displays the configuration of all Ethernet switches
(cluster and storage).
```

```
[[ID0b63ec30bd0c96690862cb9a1fabfa0d]]
= system switch ethernet show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the configuration for Ethernet switches (cluster and storage).

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet show` command displays configuration details
for cluster and storage Ethernet switches.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that have the specified name.

| ``[-banner  ]``::
Displays the following information about a switch:
+
  * Device Name
  * Banner seen when accessing the switch through SSH.



| ``[-snmp-config  ]``::
Displays the following information about a switch:
+
  * Device Name
  * SNMPv2c Community String or SNMPv3 Username
  * SNMP Version



| ``[-status  ]``::
Displays the following status information about a switch:
+
  * Is Discovered
  * SNMPv2c Community String or SNMPv3 Username
```

```
  * Model Number
  * Switch Network
  * Software Version
  * Reason For Not Monitoring
  * Source Of Switch Version
  * Is Monitored ?



| ``[-instance  ]`` }::
Selects detailed information for all the switches.

``[-device <text>]`` - Device Name::
Selects the switches that match the specified device name.

``[-address <IP Address>]`` - IP Address::
Selects the switches that match the specified IP address.

``[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}]`` - SNMP Version::
Selects the switches that match the specified SNMP version.

``[-is-discovered {true|false}]`` - Is Discovered::
Selects the switches that match the specified discovery setting.

``[-community-or-username <text>]`` - SNMPv2c Community String or SNMPv3
Username::
Selects the switches that match the specified SNMPv2c community string or
SNMPv3 username.

``[-model {OTHER|CN1610|NX3132V|NX3232C}]`` - Model Number::
Selects the switches that match the specified model number.

``[-type {cluster-network|storage-network}]`` - Switch Network::
Selects the switches that match the specified switch type.

``[-sw-version <text>]`` - Software Version::
Selects the switches that match the specified software version.

``[-reason <text>]`` - Reason For Not Monitoring::
Selects the switches that match the specified reason.

``[-version-source <text>]`` - Source Of Switch Version::
Selects the switches that match the specified version source (for example,
SNMP, CDP, ISDP, or LLDP).

``[-is-monitoring-enabled-admin {true|false}]`` - Enable Switch
Monitoring::
```

```
Selects the switches that match the specified admin monitoring status.

``[-is-monitoring-enabled-operational {true|false}]`` - Is Monitored ?::
Selects the switches that match the specified operational monitoring
status.

``[-serial-number <text>]`` - Serial Number of the Device::
Selects the switches that match the specified serial number.

``[-device-banner <text>]`` - Device Banner::
Banner seen when accessing the switch through SSH.

``[-rcf-version <text>]`` - Reference Config File Version::
Selects the switches that match the specified reference configuration file
version.

== Examples
```

cluster1::> system switch ethernet show
Switch Type Address Model
--- --- --- ---
cn1610-143—234 cluster-network 10.238.143.234 CN1610
Serial Number: 20211200007
Is Monitored: true
Reason:
Software Version: 1.1.0.1
Version Source: ISDP
cn1610-143—231 cluster-network 10.238.143.231 CN1610
Serial Number: 20211200002
Is Monitored: true
Reason:
Software Version: 1.1.0.1
Version Source: ISDP

```
The example above displays the configuration of all Ethernet switches
(cluster and storage).
```

cluster1::> system switch ethernet show -snmp-config
SNMPv2c Community
Switch or SNMPv3 Username SNMP Version
--- --- ---
SwitchA public SNMPv2c

```
The example above displays the SNMPv2c community string or SNMPv3 username
and SNMP version for all Ethernet switches (cluster and storage).
```

cluster1::> system switch ethernet show -banner

Device: SwitchA
---

*
*** NetApp Reference Configuration File (RCF)**
*** Switch : SwitchModel**
*** Filename : SwitchType-RCF-v1.8-Cluster**
*** Release Date : Apr-04-2022**
*** Version : v1.8**
*

```
The example above shows the SSH banner for all the Ethernet switches
(cluster and storage).




[[IDd7a07e53a2283acb04ad06a2d9b305cb]]
= system switch ethernet fan show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display fan information for Ethernet switches (cluster and storage).

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet fan show` command displays the status of fans
on the monitored Ethernet switches.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the specified fields.

| ``[-instance  ]`` }::
Displays detailed information for all fans on all switches.

``[-device <text>]`` - Switch Name::
Selects the fans that belong to the specified switch.

``[-fan-name <text>]`` - Fan or Fan Tray Name::
```

```
Selects the fans that match the specified fan name or fan tray name.

``[-index <integer>]`` - Sensor Index::
Selects the fans that match the specified sensor index.

``[-fan-type {single|tray}]`` - Single Fan or Fan Tray::
Selects the fans that match the specified fan type.

``[-fan-status {operational|failed|not-operational|not-present|unknown}]``
- Fan Status::
Selects the fans that match the specified operational status.

``[-display-name <text>]`` - Fan Display Name::
Selects the fans that match the specified display name.

``[-unique-name <text>]`` - Fan Unique Name::
Selects the fan that matches the specified unique name.

``[-container-name <text>]`` - Fan Container Name::
Selects the fans that match the specified container name.

``[-is-psu-fan {yes|no}]`` - Is Power Supply Unit Fan::
Selects the fans that are PSU fans (`yes` ) or are not PSU fans (`no` ).

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor::
Selects the fans that the specified health monitor continuously monitors.

``[-error-description <text>]`` - Error Description::
Selects the fans that match the specified error description.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Resource Status::
Selects the fans that match the specified status.

== Examples
```

cluster1::> system switch ethernet fan show

Switch: SwitchA
Is PSU
Fan Type Fan Status Container Fan Error
---------- ------ ---------- ---------- ------ ----------
Fan Module-1 tray operational FanBay-1 no

Fan Module-2 tray operational FanBay-2 no

Switch: SwitchB
Is PSU
Fan Type Fan Status Container Fan Error
---------- ------ ---------- ---------- ------ ----------
Fan Module-1 tray operational FanBay-1 no

Fan Module-2 tray operational FanBay-2 no

```
 The above example displays the fans and their status on the switches
names SwitchA and SwitchB.




[[ID5289807affab86ca77e2e754357dd9fd]]
= system switch ethernet interface show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display interface information for Ethernet switches (cluster and storage).

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet interface show` command displays the status
and configuration of network interfaces on the monitored switches.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields that you specify.

| ``[-counters  ]``::
Displays the current status of the following network counters:
+
  * in-octets
  * in-errors
  * in-discards
  * out-octets
  * out-errors
  * out-discards
```

```
| ``[-instance  ]`` }::
Displays detailed interface configuration for all monitored Ethernet
switches

``[-device <text>]`` - Switch Name::
Selects the interface ports that belong to the specified switch.

``[-interface-name <text>]`` - Interface Name::
Selects the interface ports that match the specified interface name.

``[-index <integer>]`` - Interface Index::
Selects the interface ports that match the specified interface index.

``[-type <interface type>]`` - Interface Type::
Selects the interface ports that match the specified interface type.

``[-mtu <integer>]`` - MTU::
Selects the interface ports that match the specified maximum transfer
unit.

``[-duplex-type {unknown|half-duplex|full-duplex}]`` - Duplex Settings::
Selects the interface ports that match the specified duplex setting.

``[-speed <integer>]`` - Interface Speed(Mbps)::
Selects the interface ports that match the specified interface speed in
bits per second.

``[-admin-state {up|down|testing}]`` - Adminstrative Status::
Selects the interface ports that match the specified administrative status
of the switch interface.

``[-oper-state {up|down|testing|unknown|dormant|not-present|lower-layer-
down}]`` - Operational Status::
Selects the interface ports that match the specified operational status.

``[-is-isl {yes|no}]`` - Is ISL::
Selects the interface ports that are Inter-Switch links (yes) or are not
Inter-Switch links (no).

``[-in-octets <Counter>]`` - Input Octets::
Selects the interface ports that match the specified number of octets
entering the interface.

``[-in-errors <Counter>]`` - Input Errors::
Selects the interface ports that match the specified number of input
```

```
packets that were dropped due to errors.

``[-in-discards <Counter>]`` - Input Discards::
Selects the interface ports that match the specified number of input
packets that were silently discarded (possibly due to buffer overflow).

``[-out-octets <Counter>]`` - Output Octets::
Selects the interface ports that match the specified number of octets that
exited the interface.

``[-out-errors <Counter>]`` - Output Errors::
Selects the interface ports that match the specified number of output
packets that were dropped due to errors.

``[-out-discards <Counter>]`` - Output Discards::
Selects the interface ports that match the specified number of output
packets that were silently discarded (possibly due to buffer overflow).

``[-interface-number <integer>]`` - Interface Number::
Selects the interface ports that match the specified interface number.

``[-unique-name <text>]`` - Interface Unique Name::
Selects the interface port that matches the specified unique name.

``[-display-name <text>]`` - Interface Display Name::
Selects the interface ports that match the specified display name.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` - Port
Config Status::
Selects the interface ports that match the specified status.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor::
Selects the interface ports that the specified health monitor continuously
monitors.

``[-switch-type {cluster-network|storage-network}]`` - Switch Network::
Selects the interface ports that match the specified switch type.

``[-remote-port-mtu <integer>,...]`` - MTU of Remote Port::
Selects the interface ports that match the specified remote port mtu.

``[-remote-port-name <text>]`` - Remote Port Name::
Selects the interface ports that match the specified remote port name.

``[-remote-device <text>]`` - Remote Device Name::
```

```
Selects the interface ports that match the specified remote device.

``[-model <text>]`` - Switch Model::
Selects the interface ports that match the specified model.

``[-mac-address <text>]`` - MAC Address::
Selects the interface ports that match the specified mac address.

``[-vlan-id <text>]`` - Vlan ID::
Selects the interface ports that match the specified vlan id.

``[-if-alias <text>]`` - Interface Alias::
Selects the interface ports that match the specified interface alias.

``[-is-vpc-peer-link {yes|no}]`` - Is vPC Peer Link Interface::
Selects the interface ports that are part of the vPC peer link (yes) or
are not part of the vPC peer link (no).

== Examples
```

cluster1::> system switch ethernet interface show

Switch: SwitchA
Opera-
Num- Admin tional Is
Interface ber Index Type Status Status ISL MTU Duplex Speed
---------- ---- --------- --------- ------- -------- ---- ----- ------ --------
Ethernet1/ 1 436207616 ethernetC up up no 1500 full- 40000
1 smacd duplex

Ethernet1/ 10 436244480 ethernetC up down no 1500 full- 40000
10 smacd duplex

Switch: SwitchB
Opera-
Num- Admin tional Is
Interface ber Index Type Status Status ISL MTU Duplex Speed
---------- ---- --------- --------- ------- -------- ---- ----- ------ --------
Ethernet1/ 1 436207616 ethernetC up up no 1500 full- 40000
1 smacd duplex

Ethernet1/ 10 436244480 ethernetC up down no 1500 full- 40000
10 smacd duplex

```
The example above displays the interfaces on all Ethernet switches
(cluster, management and storage).
```

cluster1::> system switch ethernet interface show -counters

Switch: SwitchA
In Out Out
Interface In Octets In Errors Discards Out Octets Errors Discards
--- --- --- --- --- --- ---
Ethernet1/1 1856922869 177091 0 3122212606 0 2

Ethernet1/ 3242386021 0 0 1408092011 0 74
10

Switch: SwitchB
In Out Out
Interface In Octets In Errors Discards Out Octets Errors Discards
--- --- --- --- --- --- ---
Ethernet1/1 1281177979 182012 0 3271353786 0 1

Ethernet1/ 3611218526 0 0 2626671058 0 0
10

```
The example above displays the counters on switch network interfaces for
all the Ethernet switches (cluster, management and storage).




[[IDbad6dde703a5f23dac034934bcf71182]]
= system switch ethernet log collect-support-log
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Initiates the support log collection for the specified Ethernet switch.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet log collect-support-log` command initiates the
support log collection for the specified Ethernet switch via SSH.

NOTE: Log files are removed during takeover or giveback operations.



== Parameters

``-device <text>`` - Switch Name::
Ethernet switch device for which the support log collection will be
performed.
+
NOTE: The device must be one of the devices listed as an Ethernet switch
from the xref:{relative_path}system-switch-ethernet-show.html[system
switch ethernet show] command. The full device name from the
xref:{relative_path}system-switch-ethernet-show.html[system switch
ethernet show] command must be used.
+


== Examples
```

cluster1::> system switch ethernet log collect-support-log -device cluster-sw1

```
Initiates support log collection for the specified Ethernet switches.

== Related Links

* xref:{relative_path}system-switch-ethernet-show.html[system switch
ethernet show]




[[ID688708cdae31eb33364b340fa570bd19]]
= system switch ethernet log modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the Ethernet switch log request.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet log modify` command modifies the the
enablement and status of the hourly collection of periodic data and the
status of the collection of detailed support logs.

== Parameters

``-device <text>`` - Switch Name::
Ethernet switch device for which the log request is being made.
+
NOTE: The device must be one of the devices listed as an Ethernet switch
from the xref:{relative_path}system-switch-ethernet-show.html[system
switch ethernet show] command. The full device name from the
xref:{relative_path}system-switch-ethernet-show.html[system switch
ethernet show] command must be used.
+


``[-support-requested {true|false}]`` - Support Log Requested::
Request (true) or cancel (false) support log collection.

``[-periodic-enabled {true|false}]`` - Periodic Log Enabled::
Enable (true) or disable (false) periodic log collection.
```

```
== Examples
```

cluster1::> system switch ethernet log modify -device switch-name01(Switch---SN) -periodic-enabled true

```
Enables periodic log collection for the specified Ethernet switch.

== Related Links

* xref:{relative_path}system-switch-ethernet-show.html[system switch
ethernet show]




[[IDc231411b8ad540fd0aadef2656040a93]]
= system switch ethernet log setup-password
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Obtain Ethernet switch admin passwords.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet log setup-password` command allows the
administrator to set up Ethernet switch health monitor access to certain
Ethernet switches, so that switch logs can be collected.

== Examples
```

cluster1::> system switch ethernet log setup-password
Enter the switch name: (use full name from system switch ethernet show)
Enter the password: (Enter admin password of switch)
Enter the password again: (Enter admin password of switch)
cluster1::>

```
Enables setup of switch log collection for the specified Ethernet switch.
```

```
[[ID5dc8a13b080b83fbf0c66312cd03f02]]
= system switch ethernet log show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Ethernet switch log information.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet log show` command displays the enablement and
status of the hourly collection of periodic data and the status of the
collection of detailed support logs.

== 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 this parameter is specified, the command displays information about all
entries.

``[-device <text>]`` - Switch Name::
If this parameter is specified, the command only displays log collection
details of the provided Ethernet switch(es).

``[-support-requested {true|false}]`` - Support Log Requested::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with support log collection requested
(``_true_`` ) or not requested (``_false_`` ).

``[-support-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Last Support Log
Timestamp::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided timestamp for
periodic log collection.
```

```
``[-support-filename <text>]`` - Support Log Filename::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided path for support log
collection.

``[-periodic-enabled {true|false}]`` - Periodic Log Enabled::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with periodic log collection enabled
(``_true_`` ) or disable (``_false_`` ).

``[-support-state {never-run|initiated|in-progress|completed|failed}]`` -
Support Log State::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided state; ``_never-
run_`` , ``_initiated_`` , ``_in-progress_`` , ``_completed_`` , or
``_failed_``

``[-support-status-msg <text>]`` - Support Log Status::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided status message for
support log collection.

``[-periodic-state {never-run|scheduled|in-progress|halted|failed}]`` -
Periodic Log State::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided state; either
``_never-run_`` , ``_scheduled_`` , ``_in-progress_`` , ``_halted_`` , or
``_failed_`` .

``[-periodic-status-msg <text>]`` - Periodic Log Status::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided status message for
periodic log collection.

``[-periodic-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Last Periodic Log
Timestamp::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided timestamp for
periodic log collection.

``[-periodic-filename <text>]`` - Periodic Log Filename::
If this parameter is specified, the command only displays log collection
details of the Ethernet switch(es) with the provided path for periodic log
collection.

== Examples
```

cluster1::> system switch ethernet log show
Periodic Periodic Support
Switch Log Enabled Log State Log State
--- --- --- ---
switch-name01(Switch---SN) true scheduled failed
switch-name02(Switch---SN) false failed -

```
Displays the Ethernet switches, their periodic log collection enablement,
their periodic log collection state, and their support log collection
state.




[[IDc918d1df6a1d969f6a24b194ccfd9d38]]
= system switch ethernet polling-interval modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the polling interval for Ethernet switch health

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system switch ethernet polling-interval modify` command modifies the
interval in which the Ethernet switch health monitor polls cluster,
management and storage switches.

== Parameters

``[-polling-interval <integer>]`` - Polling Interval::
Specifies the interval in which the health monitor polls switches. The
interval is in minutes. The default value is 5. The allowed range of
values is 2 to 120.

== Examples
```

cluster1::> system switch ethernet polling-interval modify -polling-interval 41

```
 Modifies the polling interval of the switches.




[[IDa21a38d0dc171f5de5eaf405c556ac9d]]
= system switch ethernet polling-interval show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the polling interval for monitoring Ethernet switch health

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet polling-interval show` command displays the
polling interval used by the Ethernet switch health monitor.

== Examples
```

cluster1::> system switch ethernet polling-interval show
Polling Interval (in minutes): 40

```
 The example above displays the polling interval period for the switches.




[[IDf431909ff7ee6c65f24f02709ebc36e6]]
= system switch ethernet power show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display power information for Ethernet switches (cluster and storage).

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description

The `system switch ethernet power show` command displays the power supply
status of the monitored Ethernet switches.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields with the names that you specify.

| ``[-instance  ]`` }::
Displays detailed power supply information for all the switches.

``[-device <text>]`` - Switch Name::
Selects the PSUs that belong to the specified switch.

``[-psu-name <text>]`` - Power Supply Name::
Selects the PSUs that match the specified power supply name.

``[-oper-status {operational|failed|not-operational|not-
present|unknown}]`` - Operational Status::
Selects the PSUs that match the specified operational status.

``[-error-description <text>]`` - Error Description::
Selects the PSUs that match the specified error description.

``[-display-name <text>]`` - Power Supply Display Name::
Selects the PSUs that match the specified display name.

``[-unique-name <text>]`` - Power Supply Unique Name::
Selects the PSU that matches the specified unique name.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` - Power
Supply Resource Health::
Selects the PSUs that match the specified status.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor::
Selects the PSUs that match the specified monitor type.

``[-admin-status {on|off|not-defined|unknown}]`` - Adminstrative Status::
Selects the PSUs that match the specified administrative status for the
power supply.

== Examples
```

```
cluster1::> system switch ethernet power show
Admin Operational
Switch Power Supply Status Status Error
----------------- ---------------- ---------- ----------- ---------------
SwitchA PowerSupply-1 on operational
SwitchA PowerSupply-2 on operational
```

```
The example above displays the power-supply status for all Ethernet
switches (cluster, management and storage).




[[ID4bc7f979101e41295cf9d41ad40da3ce]]
= system switch ethernet switch-count show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the count of cluster-network and storage-network switches.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet switch-count show` command displays the
Ethernet switch redundancy status of the cluster that the Ethernet switch
health monitor is monitoring.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>,...` parameter, the command
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
power supply information for all the switches.

``[-unique-name <text>]`` - Cluster Unique Name::
Displays the cluster name whose Ethernet switch redundancy is monitored.

``[-clus-switch-count <integer>]`` - Ethernet Switch Count::
Displays the count of cluster network switches in the cluster.
```

```
``[-mgmt-switch-count <integer>]`` - (DEPRECATED)-Management Switch
Count::
+
NOTE:  This parameter is deprecated and may be removed in a future release
of ONTAP.
+
Displays the count of management network switches in the cluster.

``[-switchless-config {true|false}]`` - 2-Node Switchless::
Displays whether the cluster is in switchless cluster configuration or
not.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor::
Displays the health monitor monitoring this cluster network switch
redundancy information of the cluster.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Ethernet Switch Redundancy Status::
Displays the Ethernet switch redundancy status of the cluster.

``[-display-name <text>]`` - Cluster Display Name::
Displays the cluster name whose Ethernet switch redundancy is monitored.

``[-stor-switch-count <integer>]`` - Storage Switch Count::
Displays the count of storage network switches in the cluster.

== Examples
```

cluster1::*> system switch ethernet switch-count show
Cluster Storage Switchless
Sw Count Sw Count Config
-------- -------- ----------
2 2 false

```
Shows the count of Ethernet switches (including cluster and storage
networks) and the switchless configuration status of the cluster.




[[ID7c0cf212f3e10f1a5e24169e9fd84067]]
= system switch ethernet temperature show
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display temperature information for Ethernet switches (cluster and
storage).

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet temperature show` command displays the
temperature status of switches monitored by the Ethernet switch health
monitor.

== 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.

``[-device <text>]`` - Switch Name::
Selects the temperature sensors that belong to the specified switch.

``[-sensor-name <text>]`` - Sensor Name::
Selects the temperature sensors that match the specified temperature
sensor name.

``[-index <integer>]`` - Sensor Index::
Selects the temperature sensors that match the specified sensor index.

``[-temperature <integer>]`` - Temperature in Celsius::
Selects the temperature sensors whose readings match the specified
temperature value.

``[-threshold-severity
{Unknown|Other|Information|Degraded|Minor|Major|Critical|Fatal}]`` -
Threshold Severity::
Selects the temperature sensors that match the specified threshold
severity.
```

```
``[-threshold-value <integer>,...]`` - Threshold Value::
Selects the temperature sensors that match the specified threshold value.

``[-sensor-status {normal|warning|alert|critical|not-present|not-
operational|unknown}]`` - Temperature Status::
Selects the temperature sensors that match the specified operational
status.

``[-display-name <text>]`` - Sensor Display Name::
Selects the temperature sensors that match the specified sensor display
name.

``[-unique-name <text>]`` - Sensor Unique Name::
Selects the temperature sensor that matches the specified unique name.

``[-monitor {node-connect|system-
connect|system|controller|chassis|cluster-switch|example|ethernet-
switch}]`` - Health Monitor::
Selects the temperature sensors that the specified health monitor
continuously monitors.

``[-error-description <text>]`` - Error Description::
Selects the temperature sensors that match the specified fault error
description.

``[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}]`` -
Resource Status::
Selects the temperature sensors that match the specified status.

== Examples
```

cluster1::> system switch ethernet temperature show
Threshold Threshold
Switch Sensor Reading Severity Value Status Error
------- ------------ ------- ---------- ---------- -------- --------------
SwitchA "Module-1, Intake-1"
24 Minor, Major
40, 50 normal
SwitchA "Module-1, Intake-2"
23 Minor, Major
40, 50 normal

```
The example above displays temperature status for all Ethernet switches
(cluster, management and storage).




[[ID786549c04cd22c62827644b9dfbd3632]]
= system switch ethernet threshold show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Ethernet switch health monitor alert thresholds

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch ethernet threshold show` command displays thresholds
used by health monitor alerts.

== Examples
```

cluster1::> system switch ethernet threshold show
Errors/Utilization Threshold is per 0.10% values: 1 = 0.10%, 5 = 0.50%
Entity-alert Threshold is the count needed to raise entity warning alert
Errors Bandwidth
Threshold (0.1%) Threshold (0.1%) Entity-alert
In Out In Out Threshold
-------- -------- -------- -------- ----------
1 1 900 900 2

```
Displays the inbound and outbound switch interface packet error thresholds
which are set at 0.1%, and the inbound and outbound switch interface
bandwidth utilization thresholds which are set at 90%. Also displays the
threshold value for entity warning alerts. The node platform health
monitor also shares the same thresholds for monitoring cluster port packet
errors on the node.




[[ID4a2219676057331f870d56a8af733f69]]
```

```
= system switch fibre-channel add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a back-end fibre-channel switch for monitoring

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch fibre-channel add` command enables you to add fibre-
channel (FC) switches for SNMP monitoring in a MetroCluster configuration.
Front end switches should not be added for monitoring and will result in a
Monitor Status Error condition.

== Parameters

``-address <IP Address>`` - FC Switch Management IP Address::
This parameter specifies the IP address of the back-end FC switch that is
added for monitoring.

``[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}]`` - Supported SNMP Version::
This parameter specifies the SNMP version that ONTAP uses to communicate
with the back-end FC switch that is added for monitoring. The default SNMP
version is SNMPv2c.

``[-snmp-community-or-username <text>]`` - SNMPv2c Community or SNMPv3
Username::
This parameter specifies the SNMPv2c community set or SNMPv3 username on
the switch that is added for monitoring.

``[-veto-backend-fabric-check {true|false}]`` - Veto Back-end Fabric
Check? [.advaced-priv]#(privilege: advanced)#::
If specified, the `system switch fibre-channel add` command will not check
if the switch is present in the MetroCluster's back-end fabric. By
default, it does not let you add switches that are not present.

``[-blades <integer>,...]`` - Cisco Director Class Switch Blades to
Monitor::
This parameter specifies the blades to monitor on the back-end switch that
is added for monitoring. It is only applicable to director-class switches.

== Examples
```

```
The following command adds a back-end switch with IP Address 10.226.197.34
for monitoring:
```

cluster1::> system switch fibre-channel add -address 10.226.197.34 -snmp-community-or-username public
cluster1::> system switch fibre-channel show
Symbolic Is Monitor
Switch Name Vendor Model Switch WWN Monitored Status
---------- -------- ------- ---------- ---------- --------- -------
Cisco_10.226.197.34
mcc-cisco-8Gb-fab-4
Cisco DS-C9148-16P-K9
2000547fee78f088 true ok
mcc-cisco-8Gb-fab-1
mcc-cisco-8Gb-fab-1
Cisco - - false -
mcc-cisco-8Gb-fab-2
mcc-cisco-8Gb-fab-2
Cisco - - false -
mcc-cisco-8Gb-fab-3
mcc-cisco-8Gb-fab-3
Cisco - - false -
4 entries were displayed.
cluster1::>

```
The following command adds a Cisco Director Class switch for monitoring.
ONTAP uses SNMPv3 and 'snmpuser1' username to communicate with this
switch.
```

cluster1::> system switch fibre-channel add -address 10.228.56.208 -snmp-version SNMPv3 -snmp-community -or-username snmpuser1 -blades 3,4

```
[[ID08b6590b4483f4e25327eaa14bc9d20c]]
= system switch fibre-channel modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify information about a back-end fibre-channel switch's configuration

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch fibre-channel modify` command enables you to modify
certain parameters for identifying and accessing the back-end fibre-
channel (FC) end switches added for monitoring in a MetroCluster
configuration.

== Parameters

``-switch-name <text>`` - FC Switch Name::
This parameter specifies the name of the switch.

``[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}]`` - SNMP Version::
This parameter specifies the SNMP version that ONTAP uses to communicate
with the switch.

``[-switch-ipaddress <IP Address>]`` - Switch IP Address::
This parameter specifies the IP address of the switch.

``[-snmp-community-or-username <text>]`` - SNMPv2c Community or SNMPv3
Username::
This parameter specifies the SNMPv2c community set or SNMPv3 username on
the switch.

``[-blades <integer>,...]`` - Director-Class Switch Blades to Monitor::
This parameter specifies the blades to monitor on the switch. It is only
applicable to director-class switches.

== Examples

The following command modifies Cisco_10.226.197.34 switch SNMP community
to 'public':
```

cluster1::> system switch fibre-channel modify -switch-name Cisco_10.226.197.34 -switch-ipaddress 10.226.197.34 -snmp-community-or-username public
cluster1::>

```
The following command modifies the blades monitored on a director-class
switch:
```

cluster1::> system switch fibre-channel modify -switch-name Cisco_10.228.56.208 -blades 3,4
cluster1::>

```
The following command modifies Brocade 6505 switch SNMP version to SNMPv3
and SNMPv3 username to 'snmpuser1':
```

cluster1::> system switch fibre-channel modify -switch-name Brocade6505 -switch-ipaddress 10.226.197.34 -snmp-version SNMPv3 -snmp-community-or-username snmpuser1
cluster1::>

```
[[ID937e6eb68daee558e1c0e968e4d4b8af]]
= system switch fibre-channel refresh
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Refresh back-end fibre-channel switch info

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `system switch fibre-channel refresh` command triggers a refresh of
the SNMP data for the MetroCluster fibre-channel (FC) switches and FC-to-
SAS bridges. It does not do anything if the refresh is already going on.
The FC switches and FC-to-SAS bridges must have been previously added for
monitoring by using the xref:{relative_path}system-switch-fibre-channel-
add.html[system switch fibre-channel add] and xref:{relative_path}system-
bridge-add.html[system bridge add] commands, respectively.

== Examples

The following command triggers a refresh for the FC switch and FC-to-SAS
bridge data:
```

cluster1::*> system switch fibre-channel refresh
cluster1::*>

```
== Related Links

* xref:{relative_path}system-switch-fibre-channel-add.html[system switch
fibre-channel add]
* xref:{relative_path}system-bridge-add.html[system bridge add]




[[ID9c25e29302f491e18110da3f023de15f]]
= system switch fibre-channel remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a back-end fibre-channel switch from monitoring

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch fibre-channel remove` command enables you to remove
back-end fibre-channel (FC) switches that were previously added for SNMP
monitoring.

== Parameters

``-switch-name <text>`` - FC Switch Name::
This parameter specifies the name of the back-end FC switch added for
monitoring.

== Examples

The following command removes 'Cisco_10.226.197.34' switch from
monitoring:
```

cluster1::> system switch fibre-channel show
Symbolic Is Monitor
Switch Name Vendor Model Switch WWN Monitored Status
---------- -------- ------- ---------- ---------- --------- -------

Cisco_10.226.197.34
mcc-cisco-8Gb-fab-4
Cisco DS-C9148-16P-K9
2000547fee78f088 true ok
mcc-cisco-8Gb-fab-1
mcc-cisco-8Gb-fab-1
Cisco - - false -
mcc-cisco-8Gb-fab-2
mcc-cisco-8Gb-fab-2
Cisco - - false -
mcc-cisco-8Gb-fab-3
mcc-cisco-8Gb-fab-3
Cisco - - false -
4 entries were displayed.
cluster1::> system switch fibre-channel remove -switch-name Cisco_10.226.197.34
cluster1::> system switch fibre-channel show
Symbolic Is Monitor
Switch Name Vendor Model Switch WWN Monitored Status
---------- -------- ------- ---------- ---------- --------- -------
mcc-cisco-8Gb-fab-4
mcc-cisco-8Gb-fab-4
Cisco
- - false -
mcc-cisco-8Gb-fab-1
mcc-cisco-8Gb-fab-1
Cisco - - false -
mcc-cisco-8Gb-fab-2
mcc-cisco-8Gb-fab-2
Cisco - - false -
mcc-cisco-8Gb-fab-3
mcc-cisco-8Gb-fab-3
Cisco - - false -
4 entries were displayed
cluster1::>

```
[[IDd2ce2efe53a16b7522e273c83b4de60a]]
= system switch fibre-channel show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display back-end fibre-channel switch information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system switch fibre-channel show` command displays information about
```

```
all the back-end fibre-channel (FC) switches in the MetroCluster
configuration. The back-end switches must have been previously added for
monitoring using the xref:{relative_path}system-switch-fibre-channel-
add.html[system switch fibre-channel add] command. If no parameters are
specified, the default command displays the following information about
the back-end FC switches:

* Switch
* Symbolic Name
* Vendor
* Model
* Switch WWN
* Is Monitored
* Monitor Status

To display detailed profile information about a single back-end FC switch,
use the `-switch-name` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Displays the specified fields for all the back-end FC switches, in column-
style output.

| ``[-connectivity  ]``::
Displays the following details about the connectivity from the back-end FC
switch to connected entities:
+
  * Port name
  * Port operating mode
  * Port world wide name
  * Peer port world wide name
  * Peer type
  * Additional information about peer

+
Displays the following details about the connectivity from the node to the
back-end FC switch:
+
  * Node name
  * Adapter name
  * Switch port name
  * Switch port speed
  * Adapter type
```

```
| ``[-cooling  ]``::
Displays the following details about the fans and temperature sensors on
the back-end FC switch:
+
  * Fan name
  * Fan speed in rotations per minute (RPM)
  * Fan operational status
  * Temperature sensor name
  * Temperature sensor reading in Celsius (C)
  * Temperature sensor status



| ``[-error  ]``::
Displays the errors related to the back-end FC switch.

| ``[-port  ]``::
Displays the following details about the back-end FC switch ports:
+
  * Port name
  * Port world wide name
  * Port administrative status
  * Port operational status
  * Port operating mode
  * Whether SFP is present in the port
  * Port speed in gigabits per second (Gbps)
  * Port BB credit
  * Peer port world wide name



| ``[-power  ]``::
Displays the following details about the back-end FC switch power
supplies:
+
  * Power supply name
  * Power supply serial number
  * Power supply operational status



| ``[-san-config  ]``::
Displays the following details about the Virtual Storage Area Networks
(VSAN) and Zones of the back-end FC switch:
+
  * VSAN identifier
  * VSAN name
```

```
  * VSAN operational status
  * Type of load balancing configured for the VSAN
  * Where in-order-delivery set for the VSAN
  * Whether the auto power reset of the PSU is enabled
  * VAN member switch name and port
  * Zone name
  * VSAN ID of the zone
  * Zone member switch name and port
  * Zone member port id
  * Zone member port world wide name



| ``[-sfp  ]``::
Displays the following details about the back-end FC switch ports small
form-factor pluggable (SFP):
+
  * Port name
  * Type of SFP
  * SFP transmitter type
  * SFP vendor
  * SFP part number
  * SFP serial number



| ``[-stats  ]``::
Displays the following details about the back-end FC switch ports:
+
  * Port name
  * Frames received through the port (Rx Frames)
  * Frames transmitted through the port (Tx Frames)
  * Octets received through the port (Rx Octets)
  * Octets transmitted through the port (Tx Octets)
  * Port error frames



| ``[-instance  ]`` }::
Displays expanded information about all the back-end FC switches in the
system. If a back-end FC switch is specified, then this parameter displays
the same detailed information for the back-end FC switch you specify as
does the -switch-name parameter.

``[-switch-name <text>]`` - FC Switch Name::
Displays information only about the back-end FC switches that match the
name you specify.
```

```
``[-switch-wwn <text>]`` - Switch World Wide Name::
Displays information only about the back-end FC switches that match the
switch wwn you specify.

``[-switch-symbolic-name <text>]`` - Switch Symbolic Name::
Displays information only about the back-end FC switches that match the
switch symbolic name you specify.

``[-switch-fabric-name <text>]`` - Fabric Name::
Displays information only about the back-end FC switches that match the
switch fabric you specify.

``[-domain-id <integer>]`` - Switch Domain ID::
Displays information only about the back-end FC switches that match the
switch domain id you specify.

``[-switch-role {unknown|primary|subordinate}]`` - Switch Role in Fabric::
Displays information only about the back-end FC switches that match the
switch role you specify.

``[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}]`` - SNMP Version::
Displays information only about the back-end FC switches that match the
switch SNMP version you specify.

``[-switch-model <text>]`` - Switch Model::
Displays information only about the back-end FC switches that match the
switch model you specify.

``[-switch-vendor {unknown|Brocade|Cisco}]`` - Switch Vendor::
Displays information only about the back-end FC switches that match the
switch vendor you specify.

``[-fw-version <text>]`` - Switch Firmware Version::
Displays information only about the back-end FC switches that match the
switch firmware version you specify.

``[-serial-number <text>]`` - Switch Serial Number::
Displays information only about the back-end FC switches that match the
switch serial number you specify.

``[-switch-ipaddress <IP Address>]`` - Switch IP Address::
Displays information only about the back-end FC switches that match the
switch IP address you specify.

``[-switch-status {unknown|ok|error}]`` - Switch Status::
Displays information only about the back-end FC switches that match the
```

```
switch status you specify.

``[-snmp-community-or-username <text>]`` - SNMPv2c Community or SNMPv3
Username::
Displays information only about the back-end FC switches that match the
switch SNMPv2c community or SNMPv3 username you specify.

``[-profile-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS
[{+|-}hh:mm]}]`` - Switch Profile Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
profile data last successful refresh timestamp you specify.

``[-is-monitoring-enabled {true|false}]`` - Is Monitoring Enabled for
Switch::
Displays information only about the back-end FC switches that match the
switch monitoring value you specify.

``[-blades <integer>,...]`` - Director-Class Switch Blades to Monitor::
Displays information only about the back-end FC switches that match the
blade value you specify.

``[-engine-id <Hex String>]`` - Engine ID of SNMPv3 Capable Switch::
Displays information only about the back-end FC switches that match the
SNMPv3 engine-id you specify.

``[-psu-name-list <text>,...]`` - Switch Power Supply Name List::
Displays information only about the back-end FC switches that have the
power supply units with the names you specify.

``[-psu-serial-number-list <text>,...]`` - Switch Power Supply Serial
Number List::
Displays information only about the back-end FC switches that have the
power supply units with the serial numbers you specify.

``[-psu-status-list {unknown|normal|warning|faulty|not-present}]`` -
Switch Power Supply Status List::
Displays information only about the back-end FC switches that have the
power supply units with the statuses you specify.

``[-psu-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Power Supply Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
power supply unit data last successful refresh timestamp you specify.

``[-temp-sensor-name-list <text>,...]`` - Switch Temperature Sensor Name
List::
Displays information only about the back-end FC switches that have the
```

```
temperature sensors with the names you specify.

``[-temp-sensor-reading-list <integer>,...]`` - Switch Temperature Sensor
Reading (C) List::
Displays information only about the back-end FC switches that have the
temperature sensors with the readings you specify.

``[-temp-sensor-status-list {unknown|normal|warning|critical}]`` - Switch
Temperature Sensor Status List::
Displays information only about the back-end FC switches that have the
temperature sensors with the statuses you specify.

``[-temp-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Temperature Sensor Data Last Successful Refresh
Timestamp::
Displays information only about the back-end FC switches that match the
temperature sensor data last successful refresh timestamp you specify.

``[-fan-name-list <text>,...]`` - Switch Fan Name List::
Displays information only about the back-end FC switches that match the
fans with the names you specify.

``[-fan-rpm-list <integer>,...]`` - Switch Fan Speed (RPM) List::
Displays information only about the back-end FC switches that match the
fans with the RPM speeds you specify.

``[-fan-status-list {unknown|operational|failed|not-operational|not-
present}]`` - Switch Fan Operational Status List::
Displays information only about the back-end FC switches that match the
fans with the statuses you specify.

``[-fan-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Fan Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
fan data last successful refresh timestamp you specify.

``[-vsan-index-list <integer>,...]`` - Switch VSAN Index List::
Displays information only about the back-end FC switches that have the
VSANs with the indexes you specify.

``[-vsan-name-list <text>,...]`` - Switch VSAN Name List::
Displays information only about the back-end FC switches that have the
VSANs with the names you specify.

``[-vsan-oper-status-list {up|down}]`` - Switch VSAN Operational Status
List::
Displays information only about the back-end FC switches that have the
```

```
VSANs with the operational statuses you specify.

``[-vsan-load-balancing-type-list {src-id-dest-id|src-id-dest-id-ox-id}]``
- Switch VSAN Load balancing Type List::
Displays information only about the back-end FC switches that have the
VSANs with the load balancing types you specify.

``[-is-vsan-iod-list {true|false}]`` - Is In-order Delivery Set for VSAN
List::
Displays information only about the back-end FC switches that have the
VSANs with the IOD setting you specify.

``[-vsan-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch VSAN Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
VSAN data last successful refresh timestamp you specify.

``[-member-switch-name-list <text>,...]`` - Member Switch List::
Displays information only about the back-end FC switches that have the
VSANs with the member switch names you specify.

``[-member-switch-port-name-list <text>,...]`` - Member Switch Port Name
List::
Displays information only about the back-end FC switches that have the
VSANs with the member switch port names you specify.

``[-vsan-id-list <integer>,...]`` - Zone VSAN ID List::
Displays information only about the back-end FC switches that have the
VSANs with the IDs you specify.

``[-zone-name-list <text>,...]`` - Switch Zone Name List::
Displays information only about the back-end FC switches that have the
zones with the names you specify.

``[-zone-member-sw-domain-id-list <integer>,...]`` - Zone Member Switch
Port Domain ID List::
Displays information only about the back-end FC switches that have the
zones with the member switch domain ids you specify.

``[-zone-member-port-name-list <text>,...]`` - Zone Member Port List::
Displays information only about the back-end FC switches that have the
zones with the port names you specify.

``[-zone-member-port-wwn-list <text>,...]`` - Zone Member WWPN List::
Displays information only about the back-end FC switches that have the
zones with the port WWNs you specify.
```

```
``[-zone-member-port-switch-name-list <text>,...]`` - Zone Member Switch
WWN List::
Displays information only about the back-end FC switches that have the
zones with the member port hosting switch names you specify.

``[-zone-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Zone Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
zone data last successful refresh timestamp you specify.

``[-zone-member-wwn-list <text>,...]`` - Zone Member WWN List::
Displays information only about the back-end FC switches that have the
zones with the member WWNs you specify.

``[-zone-member-port-id-list <text>,...]`` - Zone Member Port ID List::
Displays information only about the back-end FC switches that have the
zones with the member port ids you specify.

``[-port-wwn-list <text>,...]`` - Switch Port World Wide Name (WWPN)
List::
Displays information only about the back-end FC switches that have the
ports with the WWNs you specify.

``[-port-name-list <text>,...]`` - Switch Port Name List::
Displays information only about the back-end FC switches that have the
ports with the names you specify.

``[-port-admin-status-list {unknown|enabled|disabled}]`` - Switch Port
Admin Status List::
Displays information only about the back-end FC switches that have the
ports with administrative statuses you specify.

``[-port-oper-status-list {unknown|online|offline}]`` - Switch Port
Operational Status List::
Displays information only about the back-end FC switches that have the
ports with operational statuses you specify.

``[-port-mode-list {unknown|auto|F-port|FL-port|E-port|TE-port|U-port|G-
port|other|EX-port|D-port|SIM-port|VE-port|AE-port|AF-port}]`` - Switch
Port Mode List::
Displays information only about the back-end FC switches that have the
ports with the operating modes you specify.

``[-port-oper-speed-list <integer>,...]`` - Switch Port Current Speed (in
Gbits/sec) List::
Displays information only about the back-end FC switches that have the
ports with the operational speeds you specify.
```

```
``[-port-bb-credit-list <integer>,...]`` - Switch Port BB Credit List::
Displays information only about the back-end FC switches that have the
ports with the BB credits you specify.

``[-port-sfp-present-list {true|false}]`` - Switch Port Is SFP Present
List::
Displays information only about the back-end FC switches that have the
ports with the SFP present values you specify.

``[-port-peer-wwpn-list <text>,...]`` - Switch Port Peer WWPN List::
Displays information only about the back-end FC switches that have the
ports with the peer port WWPNs you specify.

``[-port-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Port Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
port data last successful refresh timestamp you specify.

``[-port-stat-name-list <text>,...]`` - Switch Port Name List::
Displays information only about the back-end FC switches that have the
ports with the names you specify.

``[-port-tx-frames-list <integer>,...]`` - Switch Port Transmitted Frame
Count List::
Displays information only about the back-end FC switches that have the
ports with the transmitted frames values you specify.

``[-port-rx-frames-list <integer>,...]`` - Switch Port Received Frame
Count List::
Displays information only about the back-end FC switches that have the
ports with the received frames values you specify.

``[-port-tx-octets-list <integer>,...]`` - Switch Port Total Transmitted
Octets List::
Displays information only about the back-end FC switches that have the
ports with the transmitted octets values you specify.

``[-port-rx-octets-list <integer>,...]`` - Switch Port Total Received
Octets List::
Displays information only about the back-end FC switches that have the
ports with the received octets values you specify.

``[-port-frame-error-list <integer>,...]`` - Switch Port Frame Error Count
List::
Displays information only about the back-end FC switches that have the
```

```
ports with the error frame values you specify.

``[-port-stat-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS
[{+|-}hh:mm]}]`` - Switch Port Stat Data Last Update Timestamp::
Displays information only about the back-end FC switches that match the
port statistics data last successful refresh timestamp you specify.

``[-sfp-port-name-list <text>,...]`` - Switch Port Name List::
Displays information only about the back-end FC switches that have the
ports with the names you specify.

``[-sfp-type-list {unknown|other|gbic|embedded|glm|gbic-with-serial-
id|gbic-without-serial-id|sfp-with-serial-id|sfp-without-serial-id|xfp|x2-
short|x2-medium|x2-tall|xpak-short|xpak-medium|xpak-tall|xenpak|sfp-dw-
dm|qsfp|x2-dw-dm|gbic-not-installed|small-form-factor}]`` - Switch Port
SFP Type List::
Displays information only about the back-end FC switches that have the
ports with the SFP types you specify.

``[-sfp-tx-type-list {unknown|long-wave-laser|short-wave-laser|long-wave-
laser-cost-reduced|electrical|ten-gig-base-sr|ten-gig-base-lr|ten-gig-
base-er|ten-gig-base-lx4|ten-gig-base-sw|ten-gig-base-lw|ten-gig-base-
ew}]`` - Switch Port SFP Transmitter Type List::
Displays information only about the back-end FC switches that have the
ports with the SFP transmitter types you specify.

``[-sfp-vendor-list <text>,...]`` - Switch Port SFP Vendor List::
Displays information only about the back-end FC switches that have the
ports with the SFP vendors you specify.

``[-sfp-part-number-list <text>,...]`` - Switch Port SFP Part Number
List::
Displays information only about the back-end FC switches that have the
ports with the SFP part numbers you specify.

``[-sfp-serial-number-list <text>,...]`` - Switch Port SFP Serial Number
List::
Displays information only about the back-end FC switches that have the
ports with the SFP serial numbers you specify.

``[-sfp-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Port SFP Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
port SFP data last successful refresh timestamp you specify.

``[-switch-error-text-list <text>,...]`` - Switch Error Text List::
```

```
Displays information only about the back-end FC switches that have the
errors you specify.

``[-conn-switch-port-name-list <text>,...]`` - Switch Port Name List::
Displays information only about the back-end FC switches that have the
ports with the names you specify.

``[-conn-switch-port-mode-list {unknown|auto|F-port|FL-port|E-port|TE-
port|U-port|G-port|other|EX-port|D-port|SIM-port|VE-port|AE-port|AF-
port}]`` - Switch Port Operating Mode List::
Displays information only about the back-end FC switches that have the
ports with the operating modes you specify.

``[-conn-switch-port-wwn-list <text>,...]`` - Switch Port WWN List::
Displays information only about the back-end FC switches that have the
ports with the WWNs you specify.

``[-conn-switch-port-peer-port-wwn-list <text>,...]`` - Switch Port Peer
Port WWN List::
Displays information only about the back-end FC switches that have the
ports with the peer port WWNs you specify.

``[-conn-switch-port-peer-info-list <text>,...]`` - Switch Port Peer Host
& Port Name List::
Displays information only about the back-end FC switches that have the
ports with the peer information values you specify.

``[-conn-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Connectivity Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
switch connectivity data last successful refresh timestamp you specify.

``[-conn-switch-port-peer-type-list {unknown|bridge|switch|fcp-
adapter|fcvi-adapter}]`` - Switch Port Peer Type List::
Displays information only about the back-end FC switches that have the
ports connected to the peer types you specify.

``[-switch-port-name-list <text>,...]`` - Switch Port Name List::
Displays information only about the back-end FC switches that have the
ports with the names you specify.

``[-switch-port-speed-list <integer>,...]`` - Switch Port Speed (in Gbps)
List::
Displays information only about the back-end FC switches that have the
ports with the speeds you specify.

``[-node-name-list <nodename>,...]`` - Node Name List::
```

```
Displays information only about the back-end FC switches that are
connected to the nodes you specify.

``[-adapter-name-list <text>,...]`` - Node Adapter Name List::
Displays information only about the back-end FC switches that are
connected to the adapters you specify.

``[-adapter-port-name-list <text>,...]`` - Node Adapter Port Name List::
Displays information only about the back-end FC switches that are
connected to the adapter ports you specify.

``[-adapter-type-list {unknown|FCP-Initiator|FC-VI|FCP-Target}]`` - Node
Adapter Type List::
Displays information only about the back-end FC switches that are
connected to the types of adapters you specify.

``[-path-data-last-successful-refresh-timestamp {MM/DD/YYYY HH:MM:SS [{+|-
}hh:mm]}]`` - Switch Path Data Last Successful Refresh Timestamp::
Displays information only about the back-end FC switches that match the
node to switch path data last successful refresh timestamp you specify.

``[-name-list <text>,...]`` - Switch Name List::
Displays information only about the back-end FC switches that match the
names you specify.

``[-domain-id-list <integer>,...]`` - Switch Domain ID List::
Displays information only about the back-end FC switches that match the
domain ids you specify.

``[-wwn-list <text>,...]`` - Switch WWN List::
Displays information only about the back-end FC switches that match the
switch WWNs you specify.

``[-role-list {unknown|primary|subordinate}]`` - Switch Role in Fabric
List::
Displays information only about the back-end FC switches that match the
switch roles you specify.

``[-address-list <IP Address>,...]`` - Switch IP Address List::
Displays information only about the back-end FC switches that match the
switch IP addresses you specify.

== Examples

The following example displays information about all back-end FC switches:
```

```
cluster::> system switch fibre-channel show
Symbolic Is Monitor
Switch Name Vendor Model Switch WWN Monitored Status
----------- -------- ------- ---------- ---------- --------- -------
Cisco_10.226.197.34
mcc-cisco-8Gb-fab-4
Cisco DS-C9148-16P-K9
2000547fee78f088 true ok
Cisco_10.226.197.35
mcc-cisco-8Gb-fab-3
Cisco DS-C9148-16P-K9
2000547fee78f0f0 true ok
Cisco_10.226.197.36
mcc-cisco-8Gb-fab-2
Cisco DS-C9148-16P-K9
2000547fee78efb0 true ok
Cisco_10.226.197.37
mcc-cisco-8Gb-fab-1
Cisco DS-C9148-16P-K9
2000547fee78f0d8 true ok
4 entries were displayed.
cluster::>
```

```
The following example displays connectivity (switch to peer and node to
switch) information about all back-end FC switches:
```

```
cluster::> system switch fibre-channel show -connectivity
Switch Name: Cisco_10.226.197.36
Switch WWN: 2000547fee78efb0
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9
Errors: -
Last Update Time: 7/31/2014 14:16:42 -04:00
Connectivity:
Port Name Port Mode Port WWN Peer Port WWN Peer Type Peer Info
--------- --------- ---------- ---------- ---------- ---------
fc1/1 F-port 2001547fee78efb0 2100001086607d34 unknown unknown
fc1/3 F-port 2003547fee78efb0 21000024ff3dd9cb unknown unknown
fc1/4 F-port 2004547fee78efb0 21000024ff3dda8d unknown unknown
fc1/5 F-port 2005547fee78efb0 500a0980009af880 unknown unknown
fc1/6 F-port 2006547fee78efb0 500a0981009af370 unknown unknown
fc1/11 TE-port 200b547fee78efb0 200b547fee78f088 switch Cisco_10.226.197.34:fc1/11
fc1/12 TE-port 200c547fee78efb0 200c547fee78f088 switch Cisco_10.226.197.34:fc1/12
fc1/13 F-port 200d547fee78efb0 2100001086609e22 unknown unknown
fc1/15 F-port 200f547fee78efb0 21000024ff3dd91b unknown unknown
fc1/16 F-port 2010547fee78efb0 21000024ff3dbef5 unknown unknown
fc1/17 F-port 2011547fee78efb0 500a0981009afda0 unknown unknown
fc1/18 F-port 2012547fee78efb0 500a0981009a9160 unknown unknown
fc1/25 F-port 2019547fee78efb0 21000010866037e8 bridge ATTO_10.226.197.17:1
fc1/27 F-port 201b547fee78efb0 21000024ff3dd9d3 fcvi-adapter dpg-mcc-3240-15-a1:fcvi_device_1
fc1/28 F-port 201c547fee78efb0 21000024ff3dbe3d fcvi-adapter dpg-mcc-3240-15-a2:fcvi_device_1
```

fc1/29 F-port 201d547fee78efb0 500a0980009ae0a0 fcp-adapter dpg-mcc-3240-15-a2:0c
fc1/30 F-port 201e547fee78efb0 500a0981009aef40 fcp-adapter dpg-mcc-3240-15-a1:0d
Last Update Time: 7/31/2014 14:26:48 -04:00
Path:
Switch
Port
Node Adapter Switch Port Speed Adapter Type
--- --- --- --- ---
dpg-mcc-3240-15-a1 0d fc1/30 4Gbps FCP-Initiator
dpg-mcc-3240-15-a1 fcvi_device_1 fc1/27 8Gbps FC-VI
dpg-mcc-3240-15-a2 0c fc1/29 4Gbps FCP-Initiator
dpg-mcc-3240-15-a2 fcvi_device_1 fc1/28 8Gbps FC-VI

```
The following command displays cooling (temperature sensors and fans)
information about all back-end FC switches:
```

cluster::> system switch fibre-channel show -cooling
Switch Name: Cisco_10.226.197.34
Switch WWN: 2000547fee78f088
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9
Errors: -
Last Update Time: 7/31/2014 14:26:58 -04:00
Fans:
Fan RPM Status
--- --- ---
Fan Module-1 - operational
Fan Module-2 operational
Fan Module-3 operational
Fan Module-4 operational
Last Update Time: 7/31/2014 14:27:10 -04:00
Temperature Sensors:
Sensor Temp © Status
--- --- ---
module-1 Outlet 27 normal
module-1 Outlet 29 normal
module-1 Intake 26 normal
module-1 Intake 28 normal

```
The following command displays the error information about all back-end FC
switches:
```

cluster::> system switch fibre-channel show -error
Switch Name: Cisco_10.226.197.34
Switch WWN: 2000547fee78f088

Cisco_10.226.197.34(2000547fee78f088): Switch is Unreachable over Management Network.
Switch Name: Cisco_10.226.197.35
Switch WWN: 2000547fee78f0f0

Cisco_10.226.197.35(2000547fee78f0f0): Switch is Unreachable over Management Network.
Switch Name: Cisco_10.226.197.36
Switch WWN: 2000547fee78efb0

Cisco_10.226.197.36(2000547fee78efb0): Switch is Unreachable over Management Network.
Switch Name: Cisco_10.226.197.37
Switch WWN: 2000547fee78f0d8

Cisco_10.226.197.37(2000547fee78f0d8): Switch is Unreachable over Management Network.
4 entries were displayed.

```
The following command displays the detailed information about all the
back-end FC switches:
```

cluster::> system switch fibre-channel show -instance
Switch Name: Cisco_10.226.197.34
Switch Domain: -
Switch Role: -
Switch WWN: 2000547fee78f088
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9
Firmware Version: 6.2(1)
Management IP: 10.226.197.34
Errors: Cisco_10.226.197.34(2000547fee78f088): Switch is Unreachable over Management Network.
Last Update Time: 7/31/2014 14:41:28 -04:00
Fabric:
Switch Name Domain WWN Role IP Address
--- --- --- --- ---
Cisco_10.226.197.34 0 2000547fee78f088 unknown 10.226.197.34
Cisco_10.226.197.36 0 2000547fee78efb0 unknown 10.226.197.36

```
The following command displays port information about all back-end FC
switches:
```

cluster::> system switch fibre-channel show -port
Switch Name: Cisco_10.226.197.34
Switch WWN: 2000547fee78f088
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9
Errors: -
Last Update Time: 7/31/2014 14:26:58 -04:00
Ports:
Admin Oper SFP Speed BB
Port Name Port WWN Status Status Port Mode Present (Gbps) Credit PeerPortWWN
--- --- --- --- --- --- --- --- ---
fc1/1 2001547fee78f088
enabled online F-port true 8 1 2100001086608b76

fc1/2 2002547fee78f088
enabled offline auto true 0 1
fc1/3 2003547fee78f088
enabled online F-port true 8 1 21000024ff48edd9
fc1/4 2004547fee78f088
enabled online F-port true 8 1 21000024ff3dd981
fc1/5 2005547fee78f088
enabled online F-port true 4 1 500a098001057f98
fc1/6 2006547fee78f088
enabled online F-port true 4 1 500a098101069778
fc1/7 2007547fee78f088
enabled offline auto true 0 1
fc1/8 2008547fee78f088
enabled offline auto true 0 1
fc1/9 2009547fee78f088
enabled offline auto true 0 1
fc1/10 200a547fee78f088
enabled offline auto true 0 32
fc1/11 200b547fee78f088
enabled offline TE-port true 8 32 200b547fee78efb0
fc1/12 200c547fee78f088
enabled offline TE-port true 8 32 200c547fee78efb0
fc1/13 200d547fee78f088
enabled online F-port true 8 32 2100001086609c2e
fc1/14 200e547fee78f088
enabled offline auto true 0 32
fc1/15 200f547fee78f088
enabled offline auto true 0 32
fc1/16 2010547fee78f088
enabled offline auto true 0 32
fc1/17 2011547fee78f088
enabled offline auto true 0 32
fc1/18 2012547fee78f088
enabled offline auto true 0 32
fc1/19 2013547fee78f088
enabled offline auto true 0 32
fc1/20 2014547fee78f088
enabled offline auto true 0 1
fc1/21 2015547fee78f088
enabled offline auto true 0 1
fc1/22 2016547fee78f088
enabled offline auto true 0 32
fc1/23 2017547fee78f088
enabled offline auto true 0 32
fc1/24 2018547fee78f088
enabled offline auto true 0 32
fc1/25 2019547fee78f088
enabled online F-port true 8 32 2100001086609c06
fc1/26 201a547fee78f088
enabled offline auto true 0 32
fc1/27 201b547fee78f088
enabled online F-port true 8 32 21000024ff48ea93
fc1/28 201c547fee78f088
enabled online F-port true 8 32 21000024ff48eacf

fc1/29 201d547fee78f088
enabled online F-port true 4 32 500a098101484340
fc1/30 201e547fee78f088
enabled online F-port true 4 32 500a09810147e700
fc1/31 201f547fee78f088
enabled offline auto true 0 32
fc1/32 2020547fee78f088
enabled offline auto true 0 1
fc1/33 2021547fee78f088
enabled offline auto true 0 1
fc1/34 2022547fee78f088
enabled offline auto true 0 32
fc1/35 2023547fee78f088
enabled offline auto true 0 32
fc1/36 2024547fee78f088
enabled offline auto true 0 32
fc1/37 2025547fee78f088
enabled offline auto true 0 32
fc1/38 2026547fee78f088
enabled offline auto true 0 32
fc1/39 2027547fee78f088
enabled offline auto true 0 32
fc1/40 2028547fee78f088
enabled offline auto true 0 32
fc1/41 2029547fee78f088
enabled offline auto true 0 32
fc1/42 202a547fee78f088
enabled offline auto true 0 32
fc1/43 202b547fee78f088
enabled offline auto true 0 32
fc1/44 202c547fee78f088
enabled offline auto true 0 32
fc1/45 202d547fee78f088
enabled offline auto true 0 32
fc1/46 202e547fee78f088
enabled offline auto true 0 32
fc1/47 202f547fee78f088
enabled offline auto true 0 32
fc1/48 2030547fee78f088
enabled offline auto true 0 0
port-channel 1
2401547fee78f088
enabled offline auto true 0 0
port-channel 2
2402547fee78f088
enabled offline auto true 0 0
port-channel 3
2403547fee78f088
enabled offline auto true 0 0
port-channel 4
2404547fee78f088
enabled offline auto true 0 0
port-channel 5
2405547fee78f088

enabled offline auto true 0 0
port-channel 6
2406547fee78f088
enabled offline auto true 0 0
port-channel 7
2407547fee78f088
enabled offline auto true 0 0
port-channel 8
2408547fee78f088
enabled offline auto true 0 0
port-channel 9
2409547fee78f088
enabled offline auto true 0 0
port-channel 10
240a547fee78f088
enabled offline auto true 0 0
port-channel 11
240b547fee78f088
enabled offline auto true 0 0
port-channel 12
240c547fee78f088
enabled offline auto true 0 0
sup-fc0 enabled online unknown true 1 0

```
The following command displays power supply unit information about all
back-end FC switches:
```

cluster::> system switch fibre-channel show -power
Switch Name: Cisco_10.226.197.34
Switch WWN: 2000547fee78f088
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9
Errors: -
Last Update Time: 7/31/2014 14:41:49 -04:00
Power Supplies:
Power Supply Serial Number Status
--- --- ---
300.00W 110v AC PAC15494TBZ normal
300.00W 110v AC PAC15494T4D normal

```
The following command displays san configuration (VSANs and Zones)
information about all back-end FC switches:
```

cluster::> system switch fibre-channel show -san-config
Switch Name: Cisco_10.226.197.34
Switch WWN: 2000547fee78f088
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9

```
Errors: -
Last Update Time: 7/31/2014 14:41:49 -04:00
VSAN Configuration:
Oper
VSAN ID Vsan Name Status Load Balancing isIOD
------- ------------------------ ------ --------------- -----
1 VSAN0001 up src-id-dest-id true
2 dpg_13_storage up src-id-dest-id-ox-id
true
3 dpg_13_fcvi down src-id-dest-id-ox-id
true
10 dpg_mcc_13_fab1_fcvi up src-id-dest-id true
20 dpg_mcc_13_fab1_storage up src-id-dest-id-ox-id
true
30 dpg_mcc_13_fab2_fcvi up src-id-dest-id true
40 VSAN0040 up src-id-dest-id true
70 dpg_mcc_14_fcvi up src-id-dest-id true
80 dpg_mcc_14_storage up src-id-dest-id-ox-id
true
110 dpg_mcc_15_fcvi up src-id-dest-id-ox-id
true
120 dpg_mcc_15_storage up src-id-dest-id-ox-id
true
4094 isolated_vsan down src-id-dest-id-ox-id
true
VSAN Membership:
VSAN ID Switch Name Switch Port Name
------- -------------------- ----------------
1 Cisco_10.226.197.34 fc1/2
1 Cisco_10.226.197.34 fc1/7
1 Cisco_10.226.197.34 fc1/8
1 Cisco_10.226.197.34 fc1/9
1 Cisco_10.226.197.34 fc1/10
1 Cisco_10.226.197.34 fc1/11
1 Cisco_10.226.197.34 fc1/12
1 Cisco_10.226.197.34 fc1/14
1 Cisco_10.226.197.34 fc1/19
1 Cisco_10.226.197.34 fc1/20
1 Cisco_10.226.197.34 fc1/21
1 Cisco_10.226.197.34 fc1/22
1 Cisco_10.226.197.34 fc1/23
1 Cisco_10.226.197.34 fc1/24
1 Cisco_10.226.197.34 fc1/31
1 Cisco_10.226.197.34 fc1/32
1 Cisco_10.226.197.34 fc1/33
1 Cisco_10.226.197.34 fc1/34
1 Cisco_10.226.197.34 fc1/35
1 Cisco_10.226.197.34 fc1/36
1 Cisco_10.226.197.34 fc1/37
1 Cisco_10.226.197.34 fc1/38
1 Cisco_10.226.197.34 fc1/39
1 Cisco_10.226.197.34 fc1/40
1 Cisco_10.226.197.34 fc1/41
1 Cisco_10.226.197.34 fc1/42
```

1 Cisco_10.226.197.34 fc1/43
1 Cisco_10.226.197.34 fc1/44
1 Cisco_10.226.197.34 fc1/45
1 Cisco_10.226.197.34 fc1/46
1 Cisco_10.226.197.34 fc1/47
1 Cisco_10.226.197.34 fc1/48
1 Cisco_10.226.197.34 port-channel 1
1 Cisco_10.226.197.34 port-channel 2
1 Cisco_10.226.197.34 port-channel 3
1 Cisco_10.226.197.34 port-channel 4
1 Cisco_10.226.197.34 port-channel 5
1 Cisco_10.226.197.34 port-channel 6
1 Cisco_10.226.197.34 port-channel 7
1 Cisco_10.226.197.34 port-channel 8
1 Cisco_10.226.197.34 port-channel 9
1 Cisco_10.226.197.34 port-channel 10
1 Cisco_10.226.197.34 port-channel 11
1 Cisco_10.226.197.34 port-channel 12
1 Cisco_10.226.197.36 fc1/2
1 Cisco_10.226.197.36 fc1/7
1 Cisco_10.226.197.36 fc1/8
1 Cisco_10.226.197.36 fc1/9
1 Cisco_10.226.197.36 fc1/10
1 Cisco_10.226.197.36 fc1/11
1 Cisco_10.226.197.36 fc1/12
1 Cisco_10.226.197.36 fc1/14
1 Cisco_10.226.197.36 fc1/19
1 Cisco_10.226.197.36 fc1/20
1 Cisco_10.226.197.36 fc1/21
1 Cisco_10.226.197.36 fc1/22
1 Cisco_10.226.197.36 fc1/23
1 Cisco_10.226.197.36 fc1/24
1 Cisco_10.226.197.36 fc1/26
1 Cisco_10.226.197.36 fc1/31
1 Cisco_10.226.197.36 fc1/32
1 Cisco_10.226.197.36 fc1/33
1 Cisco_10.226.197.36 fc1/34
1 Cisco_10.226.197.36 fc1/35
1 Cisco_10.226.197.36 fc1/36
1 Cisco_10.226.197.36 fc1/37
1 Cisco_10.226.197.36 fc1/38
1 Cisco_10.226.197.36 fc1/39
1 Cisco_10.226.197.36 fc1/40
1 Cisco_10.226.197.36 fc1/41
1 Cisco_10.226.197.36 fc1/42
1 Cisco_10.226.197.36 fc1/43
1 Cisco_10.226.197.36 fc1/44
1 Cisco_10.226.197.36 fc1/45
1 Cisco_10.226.197.36 fc1/46
1 Cisco_10.226.197.36 fc1/47
1 Cisco_10.226.197.36 fc1/48
30 Cisco_10.226.197.34 fc1/3
30 Cisco_10.226.197.34 fc1/4
30 Cisco_10.226.197.36 fc1/3

30 Cisco_10.226.197.36 fc1/4
40 Cisco_10.226.197.34 fc1/1
40 Cisco_10.226.197.34 fc1/5
40 Cisco_10.226.197.34 fc1/6
40 Cisco_10.226.197.36 fc1/1
40 Cisco_10.226.197.36 fc1/5
40 Cisco_10.226.197.36 fc1/6
70 Cisco_10.226.197.34 fc1/15
70 Cisco_10.226.197.34 fc1/16
70 Cisco_10.226.197.36 fc1/15
70 Cisco_10.226.197.36 fc1/16
80 Cisco_10.226.197.34 fc1/13
80 Cisco_10.226.197.34 fc1/17
80 Cisco_10.226.197.34 fc1/18
80 Cisco_10.226.197.36 fc1/13
80 Cisco_10.226.197.36 fc1/17
80 Cisco_10.226.197.36 fc1/18
110 Cisco_10.226.197.34 fc1/26
110 Cisco_10.226.197.34 fc1/27
110 Cisco_10.226.197.34 fc1/28
120 Cisco_10.226.197.34 fc1/25
120 Cisco_10.226.197.34 fc1/29
120 Cisco_10.226.197.34 fc1/30
120 Cisco_10.226.197.36 fc1/25
120 Cisco_10.226.197.36 fc1/29
120 Cisco_10.226.197.36 fc1/30
Last Update Time: 7/31/2014 14:45:40 -04:00
Zone Configuration:
Member Member Member
Zone Name VSAN ID Switch Name Port Name Port ID Member WWN
--------- ------- ----------- --------- ------- ----------
dpg_mcc_fcvi 30 Cisco_10.226.197.36
fc1/3 -
$default_zone$ 30 Cisco_10.226.197.36
fc1/4
dpg_mcc_storage
40 Cisco_10.226.197.36
fc1/1
$default_zone$ 40 Cisco_10.226.197.36
fc1/5
dpg_mcc_14_fcvi
70 Cisco_10.226.197.36
fc1/15
$default_zone$ 70 Cisco_10.226.197.36
fc1/16
dpg_mcc_14_storage
80 Cisco_10.226.197.34
fc1/13
$default_zone$ 80 Cisco_10.226.197.34
fc1/17
dpg_mcc_15_fcvi
110 Cisco_10.226.197.36
fc1/27
$default_zone$

110 Cisco_10.226.197.36
fc1/28
dpg_mcc_15_storage
120 Cisco_10.226.197.34
fc1/25
$default_zone$
120 Cisco_10.226.197.34
fc1/29

```
The following command displays port SFP information about all back-end FC
switches:
```

cluster::> system switch fibre-channel show -sfp
Switch Name: Cisco_10.226.197.34
Switch WWN: 2000547fee78f088
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9
Errors: -
Last Update Time: 7/31/2014 14:41:49 -04:00
SFP:
Port Name Type Tx Type Vendor Part Number Serial Number
--------- ---------- ---------- ---------- ---------- ----------
fc1/1 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629J9
fc1/2 unknown unknown
fc1/3 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629H3
fc1/4 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629QH
fc1/5 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160628EA
fc1/6 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629QT
fc1/7 unknown unknown
fc1/8 unknown unknown
fc1/9 unknown unknown
fc1/10 unknown unknown
fc1/11 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629GP

fc1/12 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS16061X71
fc1/13 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629P8
fc1/14 unknown unknown
fc1/15 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629JP
fc1/16 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160628D2
fc1/17 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629NG
fc1/18 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629R1
fc1/19 unknown unknown
fc1/20 unknown unknown
fc1/21 unknown unknown
fc1/22 unknown unknown
fc1/23 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629NC
fc1/24 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160628CX
fc1/25 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629NZ
fc1/26 unknown unknown
fc1/27 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS16061XB0
fc1/28 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS16061XA6
fc1/29 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS16061XA0

fc1/30 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS16061X9S
fc1/31 unknown unknown
fc1/32 unknown unknown
fc1/33 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS16061NL7
fc1/34 unknown unknown
fc1/35 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629M8
fc1/36 sfp-with-serial-id
short-wave-laser CISCO-FINISAR
FTLF8528P2BCV-CS
FNS160629KH
fc1/37 unknown unknown
fc1/38 unknown unknown
fc1/39 unknown unknown
fc1/40 unknown unknown
fc1/41 unknown unknown
fc1/42 unknown unknown
fc1/43 unknown unknown
fc1/44 unknown unknown
fc1/45 unknown unknown
fc1/46 unknown unknown
fc1/47 unknown unknown
fc1/48 unknown unknown
port-channel 1
unknown unknown
port-channel 2
unknown unknown
port-channel 3
unknown unknown
port-channel 4
unknown unknown
port-channel 5
unknown unknown
port-channel 6
unknown unknown
port-channel 7
unknown unknown
port-channel 8
unknown unknown
port-channel 9
unknown unknown
port-channel 10
unknown unknown
port-channel 11
unknown unknown
port-channel 12

unknown unknown
sup-fc0

```
The following command displays port statistics information about all back-
end FC switches:
```

cluster::> system switch fibre-channel show -stats
Switch Name: Cisco_10.226.197.34
Switch WWN: 2000547fee78f088
Fabric WWN: 2001547fee78efb1
Vendor: Cisco
Model: DS-C9148-16P-K9
Errors: -
Last Update Time: 7/31/2014 14:41:49 -04:00
Port Statistics:
Rx Rx Tx Tx Error
Port Name Frames Octets Frames Octets Frames
---------- ---------- ---------- ---------- ---------- ----------
fc1/1 2116207233 3710682580 3906335374 859905888 0
fc1/2 1 208 1 208 0
fc1/3 3238899002 903116292 3079548736 4014304952 0
fc1/4 1888758418 1643379900 2434821325 2997002344 0
fc1/5 3719731908 1808138824 1878240211 3421335100 0
fc1/6 2644430347 1042009564 249190625 2003353056 0
fc1/7 1 228 1 228 0
fc1/8 1 156 1 156 0
fc1/9 1 148 1 148 0
fc1/10 1 224 1 224 0
fc1/11 3617142898 4129927136 39089396 2595464620 0
fc1/12 473603889 1560909460 2797562521 2833496016 0
fc1/13 1852255936 1091902804 180309704 1769859928 0
fc1/14 1 140 1 140 0
fc1/15 4997082 3519688264 4283938 3370856432 0
fc1/16 4995287 3519577592 4282173 3370732136 0
fc1/17 55146756 178045212 1733567096 3030415436 0
fc1/18 63005788 4287094736 1726651844 2640371212 0
fc1/19 1 200 1 200 0
fc1/20 1 104 1 104 0
fc1/21 1 108 1 108 0
fc1/22 1 108 1 108 0
fc1/23 1 164 1 164 0
fc1/24 1 216 1 216 0
fc1/25 2810698819 1611009260 471527156 1900246656 0
fc1/26 1 104 1 104 0
fc1/27 4165019838 887421780 3848122102 2581891136 0
fc1/28 58607737 1015197080 101621078 3482734024 0
fc1/29 4266270960 222242144 3766674764 2400640552 0
fc1/30 3984658378 1443835508 152597387 678837848 0
fc1/31 1 220 1 220 0
fc1/32 1 120 1 120 0
fc1/33 1 132 1 132 0
fc1/34 1 144 1 144 0

fc1/35 1 160 1 160 0
fc1/36 1 104 1 104 0
fc1/37 1 148 1 148 0
fc1/38 1 184 1 184 0
fc1/39 1 160 1 160 0
fc1/40 1 136 1 136 0
fc1/41 1 196 1 196 0
fc1/42 1 128 1 128 0
fc1/43 1 168 1 168 0
fc1/44 1 212 1 212 0
fc1/45 1 136 1 136 0
fc1/46 1 224 1 224 0
fc1/47 1 104 1 104 0
fc1/48 1 104 1 104 0

```
== Related Links

* xref:{relative_path}system-switch-fibre-channel-add.html[system switch
fibre-channel add]




:leveloffset: -1


= system timeout commands

:leveloffset: +1


[[ID58d29f8d79baba7769fc12790b960b81]]
= system timeout modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set the CLI inactivity timeout value

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system timeout modify` command sets the timeout value for CLI
sessions. If there is no CLI activity during the length of the timeout
interval, the logged in user is logged out. The default value is 30
minutes. To prevent CLI sessions from timing out, specify a value of `0`
(zero).

== Parameters

``[-timeout <integer>]`` - Timeout (in minutes)::
Use this parameter to specify the timeout value, in minutes.

== Examples

The following example shows how to modify the timeout value for CLI
sessions to be 10 minutes:
```

cluster1::> system timeout modify -timeout 10

```
The following example shows how to prevent CLI sessions from timing out:
```

cluster1::> system timeout modify -timeout 0

```
[[IDc0fdb1d8ae53c8d4d0c60b08c56a49d1]]
= system timeout show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the CLI inactivity timeout value

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `system timeout show` command displays the timeout value for CLI
sessions. If there is no CLI activity during the length of the timeout
interval, the logged in user is logged out. A timeout value of 0 minutes
means that the CLI sessions never time out.

== Examples

The following example displays the timeout value for CLI sessions:
```

cluster1::> system timeout show
CLI session timeout: 15 minute(s)

```
:leveloffset: -1


:leveloffset: -1


= template commands

:leveloffset: +1


[[IDa7402cd09fa5d0327f9b72ca846955c7]]
= template copy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Copy a template

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Use this command to copy an existing template. The copied template becomes
a readwrite template and can be customized using template parameter family
of commands.

== Parameters

``-name <template name>`` - Name of the template::
This parameter specifies the name of the template.

``-destination-name <template name>`` - Destination template::
This parameter specifies the name of the destination template.

== Examples

The following example copies template1 to template2. The template2 will be
a readwrite template:
```

cluster1::> template copy -name template1 -destination-name template2

```
[[IDc52ba816035598f4312b2948376a479c]]
= template delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing template

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Use this command to delete an existing template.

== Parameters

``-name <template name>`` - Name of the template::
This parameter specifies the name of the template.

== Examples

The following example deletes a template named template1 from the cluster:
```

cluster1::> template delete -name template1

```
[[ID7671f2ea8b0a29a3fbfe3cf4b311f737]]
= template download
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Download a template

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Use this command to download a template from an external server to the
cluster.

== Parameters

``-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}`` - URI of
the template [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the URI from which the template will be
downloaded.

``[-name <template name>]`` - Name of the template [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name that will be assigned to the template in
the cluster.

== Examples

The following example downloads the template specified in the -uri
parameter value and names the template as template1:
```

cluster1::> template download -uri http://www.example.com/netapp-templates/mysample -name template1

```
The following example downloads the template specified in the -uri
parameter value:
```

cluster1::> template download -uri http://www.example.com/netapp-templates/template1

```
[[ID96b04f6316b6f2fd30606ef31832c09b]]
= template provision
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Provision Data ONTAP resources using the template

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The command `template provision` provisions ONTAP system based on the
template that is passed as an input to the template parameter. A wizard is
presented which will accept the required inputs.

== Parameters

``-name <template name>`` - Name of the template::
This parameter specifies the name of the template

``[-simulate {true|false}]`` - Simulate::
If this parameter is specified, the provisioning is just simulated and
there will be no changes done to the system.

== Examples

The following example provisions a vserver with required protocols using
template.
```

cluster-1::> template provision -name VserverEnvironmentSetup
Press Ctrl+C to abort.

* Setup of vserver *

Vserver Name: vs0
Vserver Language [C.UTF-8]:
Vserver Security Style [unix]:
Vserver IPSpace [Default]:

* Setup of network.interface *

Enter number of instances for object network.interface: 2
(1/2)LIF Protocol: nfs

(1/2)IP Addr: 1.1.1.1
(1/2)NetMask: 255.255.255.0
(1/2)Node Name: node1-vsim1
(1/2)Port: e0c
(2/2)LIF Protocol: nfs
(2/2)IP Addr: 1.1.1.1
(2/2)NetMask: 255.255.255.0
(2/2)Node Name: node1-vsim1
(2/2)Port: e0c

*** Setup of network.routes ***

Enter number of instances for object network.routes: 1
(1/1)Gateway: 1.1.1.1
*
*** Setup of access.dns ***
*
Search Domain: netapp.com
DNS IP Addresses List: 1.1.1.1

*** Setup of security.nis ***

NIS Domains: netapp.com
NIS IP Address: 1.1.1.1
*
* Setup of security *
*
LDAP Client Config: ldapconfig
LDAP Server IP: 1.1.1.1
LDAP Base DN: dc=examplebasedn

*** Setup of protocols ***

Protocols to configure: nfs
[Job 15] Configuring vserver for vs0 (100%)

```
[[IDbd68ed42dda8bc3af0bdefb6e3a7ed92]]
= template rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a template

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Use this command to rename an existing template.

== Parameters

``-name <template name>`` - Name of the template::
This parameter specifies the name of the template.

``-new-name <template name>`` - New name of the template::
This parameter specifies the template's new name.

== Examples

The following example renames a template template1 as template2:
```

cluster1::> template rename -name template1 -new-name template2

```
[[ID5aef73f5408ac0ec3625d3e49db9ef72]]
= template show-permissions
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Template Allowed and Disallowed System Objects

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
The template show-permission command shows all the system objects that are
allowed and disallowed for the current user.

== 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.

``[-name <text>]`` - Name::
 If you specify this parameter, only permissions that match the specified
name are displayed.

``[-object-name <text>]`` - Object Name::
 If you specify this parameter, only permissions that match the specified
object-name are displayed.

``[-permission <text>]`` - Permission::
 If you specify this parameter, only permissions that match the specified
permission are displayed.

``[-command-name <text>]`` - Command Name::
 If you specify this parameter, only permissions that match the specified
command-name are displayed.

== Examples

The following example shows all the the allowed and disallowed system
objects
```

cluster1::> template show-permissions
Template: VserverEnvironmentSetup
Object Name Command Name Permission
--- --- ---
access.dns vserver services name-service dns create allowed
network.interface network interface create allowed
network.routes network route create allowed
protocols.CIFS vserver cifs create allowed

```
[[ID45b102e46e81c2131dc533d6d654709f]]
= template show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display templates

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `template show` command displays information about templates. The
command output depends on the parameter or parameters specified with the
command. If no parameters are specified, then the command displays the
following information about all the templates:

* Template Name
* Permission (``_readonly_`` or ``_readwrite_`` )

To display detailed information about a single template, run the command
with the `-name` parameter. The detailed view provides all of the
information in the previous list with the following additional
information:

* Parent Template Name
* Description
* Version
* UUID of the Template



To display detailed information about all templates, run the command with
the `-instance` parameter.

You can specify additional parameters to display information that matches
only those parameters. For example, to display information only about
templates with readonly permissions, run the command with the `-permission
readonly` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
This parameter specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
```

```
If this parameter is specified, the command displays information about all
entries.

``[-name <template name>]`` - Name of the template::
If this parameter is specified, the command displays the detailed
information about the template that matches the specified name.

``[-permission <template_permission>]`` - Permission::
If this parameter is specified, the command displays information about the
template or templates that matches the specified permission.

``[-parent-template <text>]`` - Parent template name::
If this parameter is specified, the command displays information about the
template or templates that matches the specified parent template name.

``[-description <text>]`` - Description::
If this parameter is specified, the command displays information about the
template or templates that matches the specified description.

``[-version <text>]`` - Version::
If this parameter is specified, the command displays information about the
template or templates that matches the specified version.

``[-uuid <UUID>]`` - UUID of the template::
If this parameter is specified, the command displays information about the
template or templates that matches the specified uuid.

== Examples

The following example displays information about all templates in the
cluster:
```

cluster1::> template show
Template Permission
--- ---
template1 readonly
template2 readwrite

```
The following example displays detailed information about a template named
template1:
```

cluster1::> template show -name template1
Name of the Template: template1
Permission: readonly
Parent Template Name: -
Description: Template to configure Vserver Environment

Version: 1.0
UUID of the Template: c8dfeb58-b5c5-5697-a829-18d4ee0ba202

```
[[IDa9558529f7ef176c66814e9e82e4d986]]
= template upload
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Upload an existing template to an external server

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Use this command to upload an existing template on to an external server.

== Parameters

``-name <template name>`` - Name of the template [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the template.

``-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}`` - URI to
upload the template [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the URI to which the template will be uploaded.

== Examples

The following example uploads a template template1 on to an external
server specified in the uri input parameter:
```

cluster1::*> template upload -name template1 -uri http://www.example.com/mytemplates/

```
= template parameter commands

:leveloffset: +1


[[ID1505a3649791687556cf77b76a87d7fb]]
= template parameter modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the template parameters

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `template parameter modify` command can be used to modify the
following attributes of a template parameter:

* Default value of the parameter



== Parameters

``-template <template name>`` - Name of the template::
Name of the template.

``-name <text>`` - Name of the parameter::
This parameter specifies the name of the parameter.

``[-default-value <text>]`` - Default value of the parameter::
This parameter specifies the default value of the parameter. This value is
used by the xref:{relative_path}template-provision.html[template
provision] command when it provisions the system using this template.

== Examples

The following example modifies the default value of the parameter param1
in template template1 to value1:
```

cluster1::> template parameter modify -template template1 -parameter param1 -default-value value1

```
== Related Links

* xref:{relative_path}template-provision.html[template provision]




[[IDa131f7a29d39463d3af9955a3f510ef4]]
= template parameter show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display template parameters

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `template parameter show` command displays information about the
parameters of a template. The command output depends on the parameter or
parameters specified with the command. If no parameters are specified, the
command displays the following information about all the parameters of all
the templates in the system:

* Name of the template
* Name of the parameter
* Permission (``_readonly_`` or ``_readwrite_`` )
* Default Value for the parameter
* Readonly
* Description



To display detailed information about a single parameter of the template
run the command with the `-name` parameter. The detailed view provides all
of the information in the previous list with the following additional
information:

* Recommended Value for the parameter
* Maximum Length
```

```
* Range Maximum
* Range Minimum
* Allowed Values



To display detailed information about all the parameters of the template,
run the command with the `-instance` parameter

You can specify additional parameters to display information that matches
only those parameters. For example, to display information about all the
parameters of the templates with readonly permissions, run the command
with the `-permission readonly` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
This parameter specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
If this parameter is specified, the command displays information about all
entries.

``[-template <template name>]`` - Name of the template::
Name of the template.

``[-name <text>]`` - Name of the parameter::
If this parameter is specified, the command displays information about the
parameter of all the templates that matches the specified parameter name.

``[-permission <text>]`` - Template permission::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
permission.

``[-type <text>]`` - Type of the parameter::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
type.

``[-description <text>]`` - Parameter description::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
description.

``[-recommended-value <text>]`` - Recommended value of the parameter::
If this parameter is specified, the command displays information about the
```

```
parameter or parameters of all the templates that matches the specified
recommended value.

``[-default-value <text>]`` - Default value of the parameter::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
default value.

``[-readonly {yes|no}]`` - Read-only parameter::
If this parameter is specified with a value of true, then all the
parameters that cannot modified of all templates are displayed. If the
value specified is false, then all the parameters that can be modified of
all templates are displayed.

``[-max-length <integer>]`` - Maximum length::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
maximum length.

``[-range-max <integer>]`` - Maximum range::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
maximum range.

``[-range-min <integer>]`` - Minumum range::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
minimum range.

``[-allowed-vals <text>]`` - Allowed values::
If this parameter is specified, the command displays information about the
parameter or parameters of all the templates that matches the specified
allowed values.

== Examples

The following example displays information about all the parameters of all
the templates in the cluster:
```

cluster1::> template parameter show
Template: template1
Permission: readonly
Default Read
Parameter Type Value Only Description
--- --- --- --- ---
parameter1 string - No Parameter1
parameter2 IPAddress - No Parameter2

```
:leveloffset: -1


:leveloffset: -1


= volume commands

:leveloffset: +1


[[ID698c3d50f19c25df2f084c732e0b39f1]]
= volume autosize
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set/Display the autosize settings of the flexible volume.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume autosize` command allows the user to specify the maximum size
that a volume automatically grows to when it is out of space or the
minimum size that it shrinks to when the amount of used space is below a
certain threshold. If only the volume/Vserver name is specified, then the
current settings are displayed. This command is not supported on infinite
volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Use this parameter to specify the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the volume for which the user wants to set or
display the autosize configuration.

{ ``[-maximum-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Autosize::
Use this parameter to specify the maximum size to which a flexible volume
can grow. The default for flexible volumes is 120% of the volume size. If
the value of this parameter is invalidated by manually growing the volume
```

```
to a size more than the value of the maximum size, the maximum size is
reset to 120% of the volume size. The value for `-maximum-size` cannot be
set larger than the platform-dependent maximum flexible volume size. If
you specify a larger value, the value of `-maximum-size` is automatically
reset to the supported maximum without returning an error.The value for `-
maximum-size` supports relative sizing.

``[-minimum-size {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Autosize::
This parameter specifies the minimum size to which the volume can
automatically shrink. If the volume was created with the ``_grow_shrink_``
autosize mode enabled, then the default minimum size is equal to the
initial volume size. If the value of the `-minimum-size` parameter is
invalidated by a manual volume resize or is invalid when autosize is
enabled, the minimum size is reset to the volume size. The value for `-
minimum-size` supports relative sizing. This parameter is not supported on
infinite volumes.

``[-grow-threshold-percent <percent>]`` - Grow Threshold Used Space
Percentage::
This parameter specifies the used space threshold for the automatic growth
of the volume. When the volume’s used space becomes greater than this
threshold, the volume will automatically grow unless it has reached the
maximum autosize. This parameter is not supported on infinite volumes.

``[-shrink-threshold-percent <percent>]`` - Shrink Threshold Used Space
Percentage::
This parameter specifies the used space threshold for the automatic
shrinking of the volume. When the amount of used space in the volume drops
below this threshold, the volume will shrink unless it has reached the
specified minimum size. This parameter is not supported on infinite
volumes.

``[-mode {off|grow|grow_shrink}]`` - Autosize Mode::
This parameter specifies the autosize mode for the volume. The supported
autosize modes are:
+
  * ``_off_`` - The volume does not grow or shrink in size in response to
the amount of used space.
  * ``_grow_`` - The volume will automatically grow when used space in the
volume is above the grow threshold.
  * ``_grow_shrink_`` - The volume will grow or shrink in size in response
to the amount of used space.

+
By default, `-mode` is ``_off_`` for new flexible volumes, except for data
protection (DP) mirrors, for which the default value is ``_grow_shrink_``
```

```
. The ``_grow_`` and ``_grow_shrink_`` modes work together with snapshot
autodelete to automatically reclaim space when a volume is about to become
full. The volume parameter `-space-mgmt-try-first` controls the order in
which these two space reclamation policies are attempted.

| ``[-reset <true>]`` - Autosize Reset }::
Use this option to reset the values of autosize, max-autosize, min-
autosize, autosize-grow-threshold-percent, autosize-shrink-threshold-
percent and autosize-mode to their default values based on the current
size of the volume. For example, the max-autosize value is set to 120% of
the current size of the volume.

== Examples

The following example sets the autosize settings on a volume named vol1.
The maximum size to grow is 1TB and autogrow is enabled.
```

cluster1::> vol autosize vol1 -maximum-size 1t -mode grow
(volume autosize)
vol autosize: Flexible volume 'vs1:vol1' autosize settings UPDATED.

```
The following example shows the autosize settings on a volume named vol1.
The maximum size to grow is 1TB and autogrow is enabled.
```

cluster1::> vol autosize vol1
(volume autosize)
Volume autosize is currently ON for volume 'vs1:vol1'.
The volume is set to grow to a maximum of 1t.

```
[[ID28899068cc710489d7dd6b4102fd44bf]]
= volume create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume create` command creates a volume on a specified Vserver and
```

```
storage aggregates. You can optionally specify the following attributes
for the new volume:

* Size
* State (online, offline, or restricted)
* Type (read-write or data-protection)
* Export policy
* User ID
* Group ID
* Security style (All volume types: UNIX mode bits, CIFS ACLs, or mixed
NFS and CIFS permissions.)
* Default UNIX permissions for files on the volume
* Language
* Junction path
* Whether the junction path is active (advanced privilege level or higher
only)
* Whether the volume is the root volume for its Vserver (advanced
privilege level or higher only)
* Comment
* Whether autosizing is enabled for FlexVol volumes
* Maximum size for autosizing FlexVol volumes
* Minimum size for autosize
* Grow used space threshold percentage for autosize
* Shrink used space threshold percentage for autosize
* Whether autosizing is enabled for FlexVol volumes
* Current mode of operation of volume autosize
* Maximum directory size (advanced privilege level or higher only)
* Space guarantee style (none or volume)
* Space SLO type (none, thick or semi-thick)
* Snapshot policy
* Snapshot reserve percentage
* Use logical space reporting
* Use logical space enforcement
* Whether the volume create operation runs as a foreground or background
process
* Caching policy
* Encrypt
* File system analytics state
* Tiering object tags
* Key Manager Attribute
* SnapLock type
* Activity tracking state
* Cache retention priority
* Efficiency policy
* Tiering minimum cooling days
* Cloud retrieval policy
```

```
* Whether the volume's total number of files will be set to the highest
possible value
* Application IO Size
* Preserve unlink



== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located. If only one
data Vserver exists, you do not need to specify this parameter.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume that is to be created. A volume's
name must start with an alphabetic character (a to z or A to Z) and be 197
or fewer characters in length for FlexGroup volumes, and 203 or fewer
characters in length for all other volume types. Volume names must be
unique within a Vserver.

{ ``[-aggregate <aggregate name>]`` - Aggregate Name::
This specifies the storage aggregate on which the volume is to be created.
This parameter only applies to FlexVol volumes.

| ``[-aggr-list <aggregate name>,...]`` - List of Aggregates for FlexGroup
Constituents::
Specifies an array of names of aggregates to be used for FlexGroup
constituents. Each entry in the list will create a constituent on the
specified aggregate. The root constituent will always be placed on the
first aggregate in the list, unless `optimize-aggr-list` is specified as
``_true_`` . An aggregate may be specified multiple times to have multiple
constituents created on it. This parameter only applies to FlexGroup
volumes.

``[-aggr-list-multiplier <integer>]`` - Aggregate List Repeat Count::
Specifies the number of times to iterate over the aggregates listed with
the `-aggr-list` parameter when creating a FlexGroup. The aggregate list
will be repeated the specified number of times. Example:
```

-aggr-list aggr1,aggr2 -aggr-list-multiplier 2

```
will cause four constituents to be created in the order ``_aggr1_`` ,
``_aggr2_`` , ``_aggr1_`` , ``_aggr2_`` .
+
The default value is 4.
+
```

```
This parameter only applies to FlexGroup volumes.

``[-optimize-aggr-list {true|false}]`` - Have the System Optimize the
Order of the Aggregate List::
Specifies whether to create the constituents of the FlexGroup volume on
the aggregates specified in the `aggr-list` in the order they are
specified, or whether the system should optimize the ordering of the
aggregates. If this value is ``_true_`` , the system will optimize the
ordering of the aggregates specified in the `aggr-list` . If this value is
``_false_`` the order of the `aggr-list` will be unchanged. The default
value is ``_false_`` . This parameter only applies to FlexGroup volumes.

| ``[-auto-provision-as <FlexGroup>]`` - Automatically Provision as Volume
of Type::
Use this parameter to automatically select existing aggregates for
provisioning FlexGroup volumes. Note that the fastest aggregate type with
at least one aggregate on each node of the cluster will be selected. When
auto provisioning a FlexGroup volume, the size of the FlexGroup volume
should be a minimum of 800 GB per node.
+
When the UDO License is installed, this will be the default for volume
creation if neither the `aggregate` nor `aggr-list` parameters are
specified.
+
This parameter only applies to FlexGroup volumes.

``[-support-tiering {true|false}]`` - Automatically Provision FlexGroup on
FabricPools::
This parameter specifies whether or not FabricPools are selected when
provisioning a FlexGroup during the protection workflows using the `auto-
provision-as` parameter. Only FabricPool aggregates are used if this
parameter is set to true and only non FabricPool aggregates are used if
this parameter is set to false. Tiering support for a FlexGroup can be
changed by moving all of the constituents to the required aggregates. The
default value is ``_false_``
+
This parameter only applies to FlexGroup volumes created using the `-auto-
provision-as` parameter.

``[-use-mirrored-aggregates {true|false}]`` - Automatically Provision
FlexGroup on Mirrored Aggregates::
Use this parameter to specify whether mirrored aggregates are selected
when creating a FlexGroup using the `auto-provision-as` parameter. Only
mirrored aggregates are used if this parameter is set to true and only
unmirrored aggregates are used if this parameter is set to false.
Aggregate level mirroring for a FlexGroup can be changed by moving all of
```

```
the constituents to the required aggregates. The default value is
``_true_`` for a MetroCluster configuration and is ``_false_`` for a non-
MetroCluster configuration
+
This parameter only applies to FlexGroup volumes created using the `-auto-
provision-as` parameter.

``[-encryption-type {none|volume|aggregate|mixed}]`` - Encryption Type::
Use this parameter to specify the encryption-type while creating a
FlexGroup using the `auto-provision-as` parameter. If the value is
``_none_`` , the FlexGroup created will be unencrypted. If the value is
``_volume_`` , the FlexGroup created will be of type NVE (NetApp Volume
Encryption) and if the value is ``_aggregate_`` , the FlexGroup created
will be of type NAE (NetApp Aggregate Encryption).
+
This parameter only applies to FlexGroup volumes created using the `-auto-
provision-as` parameter.

``[-nodes {<nodename>|local}]`` - List of Nodes Hosting the Volume }::
Specifies an array of node names to be used for provisioning the
FlexGroup. If an array of node names is specified, only aggregates from
the specified nodes will be considered for provisioning. If no value is
specified, all nodes in the cluster will be used.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Size::
This optionally specifies the size of the volume. The size is specified as
a number followed by a unit designation: k (kilobytes), m (megabytes), g
(gigabytes), or t (terabytes). If the unit designation is not specified,
bytes are used as the unit, and the specified number is rounded up to the
nearest 4 KB. The minimum size for a FlexVol volume is 20 MB. The minimum
size for a volume guaranteed FlexGroup is 20 MB per constituent. The
minimum size for a none guaranteed FlexGroup is 200 MB per constituent.
However, the recommended size for all FlexGroup volumes is a minimum of
100 GB per constituent. For all volumes, the default size is set to the
minimum size. The volume's maximum size is limited by the platform
maximum. If the volume's guarantee is set to `volume` , the volume's
maximum size can also be limited by the available space in the hosting
aggregates. Volumes can be increased and decreased in size with the
xref:{relative_path}volume-modify.html[volume modify] command. The maximum
number of files a volume is configured for is listed under "Total Files"
when running the command xref:{relative_path}volume-show.html[volume
show]`-instance` .

``[-state {online|restricted|offline|force-online|force-offline|mixed}]``
- Volume State::
This optionally specifies the volume's state. A restricted volume does not
```

```
provide client access to data but is available for administrative
operations.
+
NOTE: The ``_mixed_`` state applies to FlexGroup volumes only and cannot
be specified as a target state.
+


``[-policy <text>]`` - Export Policy::
This optionally specifies the ID number of the export policy associated
with the volume. For information on export policies, see the documentation
for the xref:{relative_path}vserver-export-policy-create.html[vserver
export-policy create] command. FlexGroup volumes do not support policies
that allow NFSv4 protocol access.

``[-user <user name>]`` - User ID::
This optionally specifies the name or ID of the user that is set as the
owner of the volume's root.

``[-group <group name>]`` - Group ID::
This optionally specifies the name or ID of the group that is set as the
owner of the volume's root.

``[-security-style <security style>]`` - Security Style::
This optionally specifies the security style for the volume. Possible
values include ``_unix_`` (for UNIX mode bits), ``_ntfs_`` (for CIFS
ACLs), ``_mixed_`` (for mixed NFS and CIFS permissions) and ``_unified_``
(for mixed NFS and CIFS permissions with unified ACLs). Regardless of the
security style, both NFS and CIFS clients can read from and write to the
volume.

``[-unix-permissions <unix perm>]`` - UNIX Permissions::
This optionally specifies the default UNIX permissions for files on the
volume. Specify UNIX permissions either as a four-digit octal value (for
example, 0700) or in the style of the UNIX `ls` command (for example,
-rwxr-x---). For information on UNIX permissions, see the UNIX or Linux
documentation. The default setting is 0755 or ---rwxr-xr-x.

``[-junction-path <junction path>]`` - Junction Path::
This optionally specifies the volume's junction path. The junction path
name is case insensitive and must be unique within a Vserver's namespace.

``[-junction-active {true|false}]`` - Junction Active [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether the volume's junction path is active.
The default setting is `true` . If the junction path is inactive, the
```

```
volume does not appear in the Vserver's namespace. This parameter is
available only at the advanced privilege level and higher.

``[-vsroot {true|false}]`` - Vserver Root Volume [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether the volume is the root volume of its
Vserver. The default setting is `false` . If this parameter is set to
true, the default size of the newly created volume is 1GB.  This parameter
is not supported on FlexGroup volumes.

``[-comment <text>]`` - Comment::
This optionally specifies a comment for the volume.

``[-max-autosize {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Autosize::
This parameter allows the user to specify the maximum size to which a
volume can grow. The default for volumes is 120% of the volume size. If
the value of this parameter is invalidated by manually resizing the
volume, the maximum size is reset to 120% of the volume size. The value
for `-max-autosize` cannot be set larger than the platform-dependent
maximum volume size. If you specify a larger value, the value of `-max-
autosize` is automatically reset to the supported maximum without
returning an error.

``[-min-autosize {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Autosize::
This parameter specifies the minimum size to which the volume can
automatically shrink. If the volume was created with the ``_grow_shrink_``
autosize mode enabled, then the default minimum size is equal to the
initial volume size. If the value of the `-min-autosize` parameter is
invalidated by a manual volume resize, the minimum size is reset to the
volume size.

``[-autosize-grow-threshold-percent <percent>]`` - Autosize Grow Threshold
Percentage::
This parameter specifies the used space threshold for the automatic growth
of the volume. When the volume's used space becomes greater than this
threshold, the volume will automatically grow unless it has reached the
maximum autosize.

``[-autosize-shrink-threshold-percent <percent>]`` - Autosize Shrink
Threshold Percentage::
This parameter specifies the used space threshold for the automatic
shrinking of the volume. When the amount of used space in the volume drops
below this threshold, the volume will shrink unless it has reached the
specified minimum size.

``[-autosize-mode {off|grow|grow_shrink}]`` - Autosize Mode::
```

```
This parameter specifies the autosize mode for the volume. The supported
autosize modes are:
+
  * ``_off_`` - The volume will not grow or shrink in size in response to
the amount of used space.
  * ``_grow_`` - The volume will automatically grow when used space in the
volume is above the grow threshold.
  * ``_grow_shrink_`` - The volume will grow or shrink in size in response
to the amount of used space.

+
By default, `-autosize-mode` is ``_off_`` for new volumes, except for data
protection mirrors, for which the default value is ``_grow_shrink_`` . The
``_grow_`` and ``_grow_shrink_`` modes work together with snapshot
autodelete to automatically reclaim space when a volume is about to become
full. The volume parameter `-space-mgmt-try-first` controls the order in
which these two space reclamation policies are attempted.

``[-files-set-maximum {true|false}]`` - Set Total Files (for user-visible
data) to the Highest Value that the Volume can Hold [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether the volume's total number of files will
be set to the highest possible value. By default, it is ``_false_`` . If
``_true_`` , the volume's total number of files is set to the highest
value that the volume can hold.

``[-maxdir-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Directory Size
[.advaced-priv]#(privilege: advanced)#::
This optionally specifies the maximum directory size. The default maximum
directory size is model-dependent and optimized for the size of system
memory.

{ ``[-space-slo {none|thick|semi-thick}]`` - Space SLO::
This optionally specifies the Service Level Objective for space management
(the space SLO setting) for the volume. The space SLO value is used to
enforce volume settings so that sufficient space is set aside to meet the
space SLO. The default setting is `none` . There are three supported
values: ``_none_`` , ``_thick_`` and ``_semi-thick_`` .
+
  * `none` : The value of ``_none_`` does not provide any guarantee for
overwrites or enforce any restrictions. It should be used if the admin
plans to manually manage space consumption in the volume and aggregate,
and out of space errors.
  * `thick` : The value of ``_thick_`` guarantees that the hole fills and
overwrites to space-reserved files in this volume will always succeed by
reserving space. To meet this space SLO, the following volume-level
```

```
settings are automatically set and cannot be modified:
+
    * Space Guarantee: ``_volume_`` - The entire size of the volume is
preallocated in the aggregate. Changing the volume's `space-guarantee`
type is not supported.
    * Fractional Reserve: ``_100_`` - 100% of the space required for
overwrites is reserved. Changing the volume's `fractional-reserve` setting
is not supported.


  * `semi-thick` : The value of ``_semi-thick_`` is a best-effort attempt
to ensure that overwrites succeed by restricting the use of features that
share blocks and auto-deleting backups and snapshots in the volume. To
meet this space SLO, the following volume-level settings are automatically
set and cannot be modified:
+
    * Space Guarantee: ``_volume_`` - The entire size of the volume is
preallocated in the aggregate. Changing the volume's `space-guarantee`
type is not supported.
    * Fractional Reserve: ``_0_`` - No space will be reserved for
overwrites by default. However, changing the volume's `fractional-reserve`
setting is supported. Changing the setting to 100 means that 100% of the
space required for overwrites is reserved.
    * Snapshot Autodelete: ``_enabled_`` - Automatic deletion of snapshots
is enabled to reclaim space. To ensure that the overwrites can be
accommodated when the volume reaches threshold capacity, the following
volume snapshot autodelete parameters are set automatically to the
specified values and cannot be modified:
+
      * `enabled` : ``_true_``
      * `commitment` : ``_destroy_``
      * `trigger` : ``_volume_``
      * `defer-delete` : ``_none_``
      * `destroy-list` : ``_vol_clone_`` , ``_lun_clone_`` ,
``_file_clone_`` , ``_cifs_share_``




+
 In addition, with a value of ``_semi-thick_`` , the following
technologies are not supported for the volume:
+
      * File Clones with autodelete disabled: Only full file clones of
files, LUNs or NVMe namespaces that can be autodeleted can be created in
the volume. The use of `autodelete` for file clone create is required.
      * Partial File Clones: Only full file clones of files or LUNs that
```

```
can be autodeleted can be created in the volume. The use of `range` for
file clone create is not supported.
      * Volume Efficiency: Enabling volume efficiency is not supported to
allow autodeletion of snapshots.





| ``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
This optionally specifies the space guarantee style for the volume. A
value of ``_volume_`` reserves space on the aggregates for the entire
volume. A value of ``_none_`` reserves no space on the aggregates, meaning
that writes can fail if an aggregate runs out of space. The default
setting for the volumes on All Flash FAS systems is ``_none_`` , otherwise
the default setting is ``_volume_`` . The ``_file_`` setting is no longer
supported.

``[-type {RW|DP}]`` - Volume Type }::
This optionally specifies the volume's type, either read-write (RW) or
data-protection (DP). If you do not specify a value for this parameter, a
RW volume is created by default.

``[-snapdir-access {true|false}]`` - Snapshot Directory Access Enabled::
This optionally specifies whether clients have access to .snapshot
directories. The default setting is `true` .

``[-percent-snapshot-space <percent>]`` - Space Reserved for Snapshot
Copies::
This optionally specifies the amount of space that is reserved in the
volume for snapshots. The default setting is 5 percent, except for data
protection mirrors for which the default is 0 percent.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
This optionally specifies the snapshot policy for the volume. The default
is the snapshot policy for all volumes on the Vserver, as specified by the
`-snapshot-policy` parameter of the xref:{relative_path}vserver-
create.html[vserver create] and xref:{relative_path}vserver-
modify.html[vserver modify] commands. The schedules associated with the
`snapshot-policy` for a FlexGroup cannot have an interval shorter than 30
minutes.

``[-language <Language code>]`` - Language::
This optionally specifies the language encoding setting for the volume. By
default, the volume inherits the Vserver language encoding setting.
+
NOTE: You cannot modify the language encoding setting of a volume.
```

```
+


``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes. This
parameter applies only to FlexGroup volumes. For FlexVol volumes, the
command always runs in the foreground.

``[-nvfail {on|off}]`` - NVFAIL Option::
 Setting this optional parameter to true causes the volume to set the in-
nvfailed-state flag to true, if committed writes to the volume are lost
due to a failure. The in-nvfailed-state flag fences the volume from
further data access and prevents possible corruption of the application
data. Without specifying a value, this parameter is automatically set to
false.

``[-constituent-role <Constituent Roles>]`` - Constituent Volume Role::
This parameter is no longer supported.

{ ``[-qos-policy-group <text>]`` - QoS Policy Group Name::
 This optional parameter specifies which QoS policy group to apply to the
volume. This policy group defines measurable service level objectives
(SLOs) that do not adjust based on the volume allocated space or used
space. If you do not assign a policy group to a volume, the system will
not monitor and control the traffic to it.

| ``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name
}::
 This optional parameter specifies which QoS adaptive policy group to
apply to the volume. This policy group defines measurable service level
objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on
the volume allocated space or used space.

``[-caching-policy <text>]`` - Caching Policy Name::
This optionally specifies the caching policy to apply to the volume. A
caching policy defines how the system caches this volume's data in a Flash
Pool aggregate or Flash Cache modules. If a caching policy is not assigned
to this volume, the system uses ``_auto_`` as the default caching policy.
+
Both metadata and user data are eligible for caching. Metadata consists of
directories, indirect blocks and system metafiles. They are eligible for
read caching only. When a random write pattern is detected on user data,
the first such write is eligible for read caching while all subsequent
overwrites are eligible for write caching. The available caching policies
```

```
are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.
  * noread-random_write - Write caches all randomly overwritten user data
blocks. It does not do any read caching.
  * meta-random_write - Read caches all metadata and write caches randomly
overwritten user data blocks.
  * random_read_write-random_write - Read caches all metadata, randomly
read and randomly written user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read-random_write - Read caches all metadata, randomly read and
sequentially read user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read_random_write-random_write - Read caches all metadata,
randomly read, sequentially read and randomly written user data. It also
write caches randomly overwritten user data blocks.
  * all-random_write - Read caches all data blocks read and written. It
also write caches randomly overwritten user data blocks.

+
Note that in a caching-policy name, a hyphen (-) separates read and write
policies. Default caching-policy is auto.

``[-cache-retention-priority {normal|low|high}]`` - Cache Retention
Priority [.advaced-priv]#(privilege: advanced)#::
This optionally specifies the cache retention priority to apply to the
volume. A cache retention priority defines how long the blocks of a volume
will be cached in Flash Pool once they become cold. If a cache retention
priority is not assigned to this volume, the system uses the default
policy. This parameter is available only at the advanced privilege level
and higher.
+
The available cache retention priority are:
```

```
+
  * low - Cache the cold blocks for the lowest time.
  * normal - Cache the cold blocks for the default time.
  * high - Cache the cold blocks for the highest time.



``[-is-autobalance-eligible {true|false}]`` - Is Eligible for Auto Balance
Aggregate [.advaced-priv]#(privilege: advanced)#::
If the Auto Balance feature is enabled, this parameter specifies whether
the volume might be considered for system workload balancing. When set to
``_true_`` , the Auto Balance Aggregate feature might recommend moving
this volume to another aggregate. The default value is ``_true_`` .

``[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum size of
a FlexGroup Constituent [.advaced-priv]#(privilege: advanced)#::
This optionally specifies the maximum size of a FlexGroup constituent. The
default value is determined by checking the maximum FlexVol size setting
on all nodes used by the FlexGroup. The smallest value found is selected
as the default for the `-max-constituent-size` for the FlexGroup. This
parameter applies to FlexGroup volumes only.

``[-efficiency-policy <efficiency policy>]`` - Storage Efficiency Policy
[.advaced-priv]#(privilege: advanced)#::
 This optionally specifies which storage efficiency policy to apply to the
volume. This parameter is applicable only for All-Flash FAS. This
parameter is not supported on data protection volumes on any platform. To
disable all efficiency features on the volume in All-Flash FAS, use the
value ``_none_`` . The default value is ``_auto_`` . The policies
``_inline-only_`` and ``_none_`` are not supported on Capacity optimized
Flash with QAT supported platforms.

``[-snaplock-type {non-snaplock|compliance|enterprise}]`` - SnapLock
Type::
This optionally specifies the SnapLock type of the volume. The default
value is ``_non-snaplock_`` .

``[-vserver-dr-protection {protected|unprotected}]`` - Vserver DR
Protection::
 This optionally specifies whether the volume should be protected by
Vserver level SnapMirror. This parameter is applicable only if the Vserver
is the source of a Vserver level SnapMirror relationship. The default
value for a volume of type “RW” is ``_protected_`` .

``[-encrypt {true|false}]`` - Enable or Disable Encryption::
This parameter allows the user to create an encrypted volume. When it is
```

```
set to `true` , a new key is generated, and the volume is encrypted using
the generated key. When it is set to `false` , the volume created is
unencrypted.

``[-is-space-reporting-logical {true|false}]`` - Logical Space Reporting::
This optionally specifies whether to report space logically on the volume.
When space is reported logically, ONTAP reports the volume space such that
all the physical space saved by the storage efficiency features are also
reported as used. The default setting is ``_false_`` .

``[-is-space-enforcement-logical {true|false}]`` - Logical Space
Enforcement::
This optionally specifies whether to perform logical space accounting on
the volume. When space is enforced logically, ONTAP enforces volume
settings such that all the physical space saved by the storage efficiency
features will be calculated as used. The default setting is ``_false_`` .

``[-tiering-policy <Tiering Policy>]`` - Volume Tiering Policy::
This optional parameter specifies the tiering policy to apply to the
volume. This policy determines whether the user data blocks of a volume in
a FabricPool will be tiered to the cloud tier when they become cold.
FabricPool combines Flash (performance tier) with an object store (cloud
tier) into a single aggregate. The default tiering policy is ``_snapshot-
only_`` for a FlexVol and ``_none_`` for a FlexGroup. Temperature of a
volume block increases if it is accessed frequently and decreases when it
is not.
+
The available tiering policies are:
+
  * snapshot-only - This policy allows tiering of only the volume
snapshots not associated with the active file system. The default minimum
cooling period is 2 days. The `-tiering-minimum-cooling-days` parameter
can be used to override the default.
  * auto - This policy allows tiering of both snapshot data and active
file system user data to the cloud tier. The default cooling period is 31
days. The `-tiering-minimum-cooling-days` parameter can be used to
override the default.
  * none - Volume blocks will not be tiered to the cloud tier.
  * all - This policy allows tiering of both snapshot data and active file
system user data to the cloud tier as soon as possible without waiting for
a cooling period. On DP volumes, this policy allows all transferred user
data blocks to start in the cloud tier.



``[-cloud-retrieval-policy {default|on-read|never|promote}]`` - Volume
```

```
Cloud Retrieval Policy [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the cloud retrieval policy for the
volume. This policy determines which tiered out blocks to retrieve from
the capacity tier to the performance tier.
+
The available cloud retrieval policies are:
+
  * default - This policy retrieves tiered data based on the underlying
tiering policy. If the tiering policy is 'auto', tiered data is retrieved
only for random client driven data reads. If the tiering policy is 'none'
or 'snapshot-only', tiered data is retrieved for random and sequential
client driven data reads. If the tiering policy is 'all', tiered data is
not retrieved.
  * on-read - This policy retrieves tiered data for all client driven data
reads.
  * never - This policy never retrieves tiered data.
  * promote - This policy retrieves all eligible tiered data automatically
during the next scheduled scan. It is only supported when the tiering
policy is 'none' or 'snapshot-only'. If the tiering policy is 'snapshot-
only', the only data brought back is the data in the AFS. Data that is
only in a snapshot stays in the cloud.



``[-tiering-minimum-cooling-days <integer>]`` - Volume Tiering Minimum
Cooling Days [.advaced-priv]#(privilege: advanced)#::
 This optional parameter specifies the minimum number of days that user
data blocks of the volume must be cooled before they can be considered
cold and tiered out to the cloud tier. For volumes hosted on FabricPools,
this parameter is used for tiering purposes and does not affect the
reporting of inactive data. For volumes hosted on non-FabricPools, this
parameter affects the cooling window used for the reporting of inactive
data. The value specified should be greater than the frequency with which
applications in the volume shift between different sets of data. Valid
values are between 2 and 183. This parameter cannot be set when volume
tiering policy is either "none" or "all". The default value of this
parameter depends on the volume's tiering policy. See the tiering policy
section of this documentation for corresponding default values. If the
tiering policy on the volume gets changed, then this parameter will be
reset to the default value corresponding to the new tiering policy.

``[-tiering-object-tags <text>,...]`` - Tags to be Associated with Objects
Stored on a FabricPool::
This optional parameter specifies tiering object tags to be associated
with objects stored on a FabricPool.
+
```

```
Object tags should follow these rules:
+
  * Each object tag should be a key-value pair separated by '='.
  * Multiple tags should be separated by ','. Overall tags should be in
format key1=value1[,key2=value2,...].
  * All tags of a volume must have a unique key.
  * Each tag key should start with either a letter or an underscore. Keys
should contain only alphanumeric characters and underscores. Maximum
allowed limit is of 127 characters.
  * Each tag value should be of maximum 127 characters consisting of only
alphanumeric characters and underscores.
  * Maximum 4 object tags are allowed per volume.



{ ``[-application-io-size <Application I/O Size>]`` - Block Size to Use by
Compression [.advaced-priv]#(privilege: advanced)#::
This optionally specifies the application IO size to apply to the volume.
This parameter specifies the compression group size for enabling
compression on the volume. This parameter is applicable only for All Flash
FAS. This parameter is not supported on data protection volumes on any
platform.
+
The available application I/O Sizes are:
+
  * 8K
  * auto - ONTAP will select compression group size automatically, based
on temperature of the block.



``[-storage-efficiency-mode {default|efficient}]`` - Storage Efficiency
Mode::
This optionally specifies efficiency mode of the volume.
+
The available efficiency modes are:
+
  * default
  * efficient


+
The `default` mode sets 8k adaptive compression on the volume. The
`efficient` mode sets auto adaptive compression on the volume. This
parameter is supported only on AFF platform. Auto adaptive compression
will be set on QAT supported platforms irrespective of the modes.
```

```
``[-analytics-state <Analytics State>]`` - File System Analytics State::
 This optionally specifies whether the volume will keep track of file
system usage for the purpose of analytics. The default value is ``_off_``
.

``[-activity-tracking-state <Activity Tracking State>]`` - Activity
Tracking State::
 This optionally specifies whether the volume will keep track of volume
activity, in real time. Tracked activities include read and write traffic
for files, directories, users and clients. The default value is ``_off_``
.

``[-key-manager-attribute <text>]`` - Key Manager Attribute::
This optional parameter specifies an additional key manager attribute that
should be an identifier-value pair separated by '=', ex. CRN=my-unique-
value
+
The following identifiers are currently supported:
+
  * CRN



``[-anti-ransomware-state {disabled|enabled|dry-run|paused|dry-run-
paused|enable-paused|disable-in-progress}]`` - Anti-ransomware State::
Use this parameter to specify the anti-ransomware-state while creating a
volume. If the value is ``_disabled_`` , the volume created will have
anti-ransomware feature disabled. If the value is ``_enabled_`` , the
volume created will have anti-ransomware feature enabled and if the value
is ``_dry-run_`` , the volume created will be in the anti-ransomware
evaluation mode.

``[-granular-data {disabled|basic|advanced}]`` - Granular data::
This parameter specifies the type of granular data storage supported on
the volume. Once this property is changed to any value other than
``_disabled_`` on a volume, it cannot be disabled on that volume using
this parameter. It can only be disabled on that volume by restoring a
snapshot. This property is only supported on FlexGroup volumes. The
default value is ``_basic_`` for AWS S3 FlexGroup volumes and
``_disabled_`` otherwise. Files modified by the granular data algorithms
use granular data, which might impact read/write I/O performance.
+
The available options are:
+
  * disabled - No volume capacity rebalancing operations are performed.
  * basic - Volume capacity rebalancing operations are enabled.
```

```
  * advanced - All ``_basic_`` features are enabled. In addition, granular
data distribution is enabled. Granular data distribution allows for
improved capacity rebalancing and usage balancing in certain scenarios.
This property can be set when the current value for the volume is
``_disabled_`` or ``_basic_`` . Once this property is set on a volume, it
cannot be modified to ``_disabled_`` or ``_basic_`` using this parameter.
It can only be changed on that volume by restoring a snapshot.



``[-atime-update-period <integer>]`` - Access Time Update Period (Seconds)
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the time period that must be elapsed in between
consecutive atime update events. Note: The value of this parameter is only
used if `-atime-update` is enabled on the volume.

``[-snapshot-locking-enabled {true|false}]`` - Enable Snapshot Copy
Locking::
This parameter specifies whether the snapshot locking feature is enabled,
based on an independent compliance-clock time. Once this property is
enabled, it cannot be disabled and the volume cannot be deleted until all
locked snapshots are past their expiry time. The default value is false.

``[-is-large-size-enabled {true|false}]`` - Are Large Size Volumes and
Files Enabled::
This parameter specifies whether support for large Flexible volumes and
large files is enabled. If this value is true, the maximum volume size for
a Flexible volume will be 300TB and the maximum size of a single file will
be 128TB. If this value is false, the maximum volume size for a Flexible
volume will be 100TB and the maximum size of a single file will be 16TB.
The default value is false.

``[-is-preserve-unlink-enabled {true|false}]`` - Is Preserve Unlink
Enabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether preserve unlink is enabled for NFSV4.1
shares on the volume. Once this feature is enabled, a file removal op on a
file with existing share locks will result in the file being moved to the
trash directory. Clients will continue to be able to access the file using
its file handle. The last close for that file will result in deletion of
the file. The default value is false.

``[-is-cloud-write-enabled {true|false}]`` - Is Cloud Write Enabled
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether cloud write is enabled on the volume.
This feature is only available on volumes in FabricPools. The default
value is false.
```

```
``[-aggressive-readahead-mode
{none|file_prefetch|sequential_read|cross_file_sequential_read}]`` -
Aggressive readahead mode [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the aggressive readahead mode of the volume.
Possible values are ``_none_`` and ``_file_prefetch_`` . When set to
``_file_prefetch_`` , on a file read, the system aggressively issues
readaheads for all of the blocks in the file and retains those blocks in a
cache for a finite period of time. This feature is only available on
FabricPool volumes on FSx for ONTAP and Cloud Volumes ONTAP. The default
value is ``_none_`` .

``[-in-consistency-group {true|false}]`` - If this Volume is part of a
Consistency Group::
This parameter specified if this volume is associated with a consistency
group.

== Examples

The following example creates a new volume named user_jdoe on a Vserver
named vs0 and a storage aggregate named aggr1. Upon its creation, the
volume is placed in the online state. It uses the export policy named
default_expolicy. The owner of the volume's root is a user named jdoe
whose primary group is named dev. The volume's junction path is
/user/jdoe. The volume is 250 GB in size, space for the entire volume is
reserved on the aggregate, and the create operation runs in the
background.
```

cluster1::> volume create -vserver vs0 -volume user_jdoe -aggregate aggr1
-state online -policy default_expolicy -user jdoe -group dev
-junction-path /user/jdoe -size 250g -space-guarantee volume
-percent-snapshot-space 20 -foreground false

```
The following example creates a new volume named vol_cached on a Vserver
named vs0 and a Flash Pool storage aggregate named aggr1. The newly
created volume is placed online and uses auto as the caching policy.
```

cluster1::> volume create -vserver vs0 -volume vol_cached -aggregate aggr1
-state online -caching-policy auto

```
The following example creates a new FlexGroup named media_vol on a Vserver
named vs0 with four constituents on aggregates aggr1 and aggr2. Upon its
creation, the volume is placed in the online state. The volume's junction
path is /media. The volume is 200 TB in size, no space for the volume is
reserved on the aggregates, and the create operation runs in the
background.
```

cluster1::> volume create -vserver vs0 -volume media_vol
-aggr-list aggr1,aggr1,aggr2,aggr2 -junction-path /media -size 200TB
-space-guarantee none -foreground false

```
The following example creates a new FlexGroup volume named fg on a Vserver
named vs0 on aggregates selected by ONTAP.
```

cluster1::> volume create -vserver vs0 -volume fg -auto-provision-as flexgroup

```
== Related Links

* xref:{relative_path}volume-modify.html[volume modify]
* xref:{relative_path}volume-show.html[volume show]
* xref:{relative_path}vserver-export-policy-create.html[vserver export-
policy create]
* xref:{relative_path}vserver-create.html[vserver create]
* xref:{relative_path}vserver-modify.html[vserver modify]




[[ID2e07720531b1646c2e40ecc393affa89]]
= volume delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume delete` command deletes the specified volumes. Before deleting
```

```
a volume, the user is prompted to confirm the operation unless the `-
force` flag is specified. If this volume was associated with a policy
group the underlying qos workload is deleted.

NOTE:

* If there is a qtree or quota policy associated with a volume, it is
deleted when you delete the volume.
* A volume must be offline (see xref:{relative_path}volume-
offline.html[volume offline] ) to be deleted.
*





== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver from which the volume is to be
deleted. If only one data Vserver exists, you do not need to specify this
parameter.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume that is to be deleted.

``[-force <true>]`` - Force Delete [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the user is not prompted to confirm each
deletion operation. In addition, the operation is run only on the local
node, and several potential errors are ignored. By default, this setting
is `false` . This parameter is available only at the advanced privilege
level and higher.

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes.

== Examples

The following example deletes a volume named vol1_old from a Vserver named
vs0:
```

cluster1::> volume delete -vserver vs0 -volume vol1_old

```
== Related Links

* xref:{relative_path}volume-offline.html[volume offline]




[[ID2a823a942601a2606c49cfbbf8b9565e]]
= volume expand
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Expand the size of a volume by adding constituents

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume expand` command allows the user to increase the size of a
FlexGroup by adding constituents. The size of the new constituents is
determined by the size of the smallest existing constituent. This command
only applies to FlexGroup volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter can be used to specify the Vserver on which the volume is
located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the volume for which the user wants to expand.

``-aggr-list <aggregate name>,...`` - List of Aggregates for FlexGroup
Constituents::
Specifies an array of names of aggregates to be used for new FlexGroup
constituents. Each entry in the list will create a constituent on the
specified aggregate. An aggregate may be specified multiple times to have
multiple constituents created on it.

``[-aggr-list-multiplier <integer>]`` - Aggregate List Repeat Count::
Specifies the number of times to iterate over the aggregates listed with
the `-aggr-list` parameter when expanding a FlexGroup. The aggregate list
will be repeated the specified number of times. Example:
```

-aggr-list aggr1,aggr2 -aggr-list-multiplier 2

```
will cause four constituents to be created in the order ``_aggr1_`` ,
``_aggr2_`` , ``_aggr1_`` , ``_aggr2_`` . The default value is 1.

``[-foreground {true|false}]`` - Foreground Process::
If ``_false_`` is specified for this parameter, the command runs as a job
in the background. If ``_true_`` is specified, the command will not return
until the operation is complete. The default value is ``_true_`` .

== Examples

The following example increases the size of a FlexGroup by adding 3
constituents:
```

cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume size
------- --------- -----
vs1 flexgroup 180TB

```
      cluster1::> volume expand -vserver vs1 -volume flexgroup -aggr-list
aggr1,aggr2,aggr3
Warning: The following number of constituents of size 20TB will be added
to
      FlexGroup "flexgroup": 3. Expanding the FlexGroup will cause the
state of
      all snapshots to be set to "partial". Partial snapshots
      cannot be restored.
      Do you want to continue? {y|n}: y
      [Job 52] Job succeeded: Successful
```

```
cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume    size
------- --------- -----
vs1     flexgroup 240TB
```

```
The following example increase the size of a FlexGroup by adding 6
constituents using the `-aggr-list-multiplier` :
```

cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume size
------- --------- -----
vs1 flexgroup 240TB

```
      cluster1::> volume expand -vserver vs1 -volume flexgroup -aggr-list
aggr1,aggr2 -aggr-list-multiplier 3
Warning: The following number of constituents of size 20TB will be added
to
      FlexGroup "flexgroup": 6. Expanding the FlexGroup will cause the
state of
      all snapshots to be set to "partial". Partial snapshots
      cannot be restored.
      Do you want to continue? {y|n}: y
      [Job 53] Job succeeded: Successful
```

```
cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume    size
------- --------- -----
vs1     flexgroup 360TB
```

```
[[ID2db6745ad10396998c7ae059f1c0f813]]
= volume make-vsroot
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Designate a non-root volume as a root volume of the Vserver

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume make-vsroot` command promotes a non-root volume of the Vserver
to be the Vserver's root volume. The Vserver's root volume must be a
FlexVol volume with a size of at least 1 GB.

For instance, if you run this command on a volume named ``_user_`` that is
located on a Vserver named ``_vs0_`` , the volume ``_user_`` is made the
root volume of the Vserver ``_vs0_`` .

This command is available only at the advanced privilege level and higher.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which a non-root volume is to be made the
root volume.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the non-root volume that is to be made the root volume of
its Vserver. This must be an existing FlexVol volume. Using a SnapLock
volume as the root volume for a Vserver is not supported.

== Examples

The following example makes a volume named ``_root_vs0_backup_`` the root
volume of its Vserver with FlexVol volumes, which is named ``_vs0_`` .
```

node::> volume make-vsroot -vserver vs0 -volume root_vs0_backup

```
The following example makes a volume named ``_root_vs1_`` the root volume
of the Vserver with Infinite Volume ``_vs1_`` .
```

node::> volume make-vsroot -vserver vs1 -volume root_vs1 -aggregate aggr1

```
[[ID27abdba860fbc6851e705d573dc2b59b]]
= volume modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify volume attributes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume modify` command can be used to modify the following attributes
of a volume:

* Size
* State (online, offline, restricted, force-online or force-offline)
* Export policy
* User ID
* Group ID
* Security style (All volume types: UNIX mode bits, CIFS ACLs, or mixed
NFS and CIFS permissions.)
* Default UNIX permissions for files on the volume
* Whether the junction path is active
* Comment
* Volume nearly full threshold percent
* Volume full threshold percent
* Maximum size for autosizing
* Minimum size for autosize
* Grow used space threshold percentage for autosize
* Shrink used space threshold percentage for autosize
* Whether autosizing is enabled
* Current mode of operation of volume autosize
* Reset the autosize values to their defaults
* Total number of files for user-visible data permitted on the volume
* Space guarantee style (none or volume)
* Space SLO type (none, thick or semi-thick)
```

```
* Snapshot policy
* Use logical space reporting
* Use logical space enforcement
* Tiering object tags
* Convert ucode
* Whether the volume's total number of files will be set to the highest
possible value
* Caching policy
* Cache retention priority
* Tiering minimum cooling days
* Cloud retrieval policy
* Preserve unlink



You can use the `volume move` command to change a volume's aggregate or
node. You can use the xref:{relative_path}volume-rename.html[volume
rename] command to change a volume's name. You can use the
xref:{relative_path}volume-make-vsroot.html[volume make-vsroot] command to
make a volume the root volume of its Vserver.

You can change additional volume attributes by using this command at the
advanced privilege level and higher.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located. If only one
data Vserver exists, you do not need to specify this parameter. Although
node Vservers are not displayed when using <Tab> completion, this
parameter supports node Vservers for modifying the root volume of the
specified node Vserver.

``-volume <volume name>`` - Volume Name::
This specifies the volume that is to be modified.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Size::
This optionally specifies the new size of the volume. The size is
specified as a number followed by a unit designation: k (kilobytes), m
(megabytes), g (gigabytes), or t (terabytes). If the unit designation is
not specified, bytes are used as the unit, and the specified number is
rounded up to the nearest 4 KB. A relative rather than absolute size
change can be specified by adding + or - before the given size: for
example, specifying +30m adds 30 megabytes to the volume's current size.
The minimum size for a volume is 20 MB (the default setting). The volume's
maximum size is limited by the platform maximum. If the volume's guarantee
```

```
is set to `volume` , the volume's maximum size can also be limited by the
available space in the hosting aggregate. If the volume's guarantee is
currently disabled, its size cannot be increased.

``[-state {online|restricted|offline|force-online|force-offline|mixed}]``
- Volume State::
This optionally specifies the volume's state. A restricted volume does not
provide client access to data but is available for administrative
operations.
+
NOTE: The ``_mixed_`` state applies to FlexGroup volumes only and cannot
be specified as a target state.
+



``[-policy <text>]`` - Export Policy::
This optionally specifies the ID number of the export policy associated
with the volume. For information on export policy, see the documentation
for the xref:{relative_path}vserver-export-policy-create.html[vserver
export-policy create] command. FlexGroup volumes do not support export
policies that allow NFSv4 protocol access.

``[-user <user name>]`` - User ID::
This optionally specifies the name or ID of the user that is set as the
owner of the volume's root.

``[-group <group name>]`` - Group ID::
This optionally specifies the name or ID of the group that is set as the
owner of the volume's root.

``[-security-style <security style>]`` - Security Style::
This optionally specifies the security style for the volume. Possible
values include ``_unix_`` (for UNIX mode bits), ``_ntfs_`` (for CIFS
ACLs), ``_mixed_`` (for mixed NFS and CIFS permissions) and ``_unified_``
(for mixed NFS and CIFS permissions with unified ACLs). Regardless of the
security style, both NFS and CIFS clients can read from and write to the
volume.

``[-unix-permissions <unix perm>]`` - UNIX Permissions::
This optionally specifies the default UNIX permissions for files on the
volume. Specify UNIX permissions either as a four-digit octal value (for
example, 0700) or in the style of the UNIX ls command (for example, -rwxr
-x---). For information on UNIX permissions, see the UNIX or Linux
documentation. The default setting is 0755 or -rwxr-xr-x.

``[-junction-active {true|false}]`` - Junction Active [.advaced-
```

```
priv]#(privilege: advanced)#::
This optionally specifies whether the volume's junction path is active.
The default setting is `true` . If the junction is inactive, the volume
does not appear in the Vserver's namespace.

``[-comment <text>]`` - Comment::
This optionally specifies a comment for the volume.

``[-space-nearly-full-threshold-percent <percent>]`` - Volume Nearly Full
Threshold Percent::
This optionally specifies the percentage at which the volume is considered
nearly full, and above which an EMS warning will be generated. The default
value is 95%. The maximum value for this option is 99%. Setting this
threshold to 0 disables the volume nearly full space alerts.

``[-space-full-threshold-percent <percent>]`` - Volume Full Threshold
Percent::
This optionally specifies the percentage at which the volume is considered
full, and above which a critical EMS error will be generated. The default
value is 98%. The maximum value for this option is 100%. Setting this
threshold to 0 disables the volume full space alerts.

{ ``[-max-autosize {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Autosize::
This parameter allows the user to specify the maximum size to which a
volume can grow. The default for volumes is 120% of the volume size. If
the value of this parameter is invalidated by manually resizing the
volume, the maximum size is reset to 120% of the volume size. The value
for `-max-autosize` cannot be set larger than the platform-dependent
maximum volume size. If you specify a larger value, the value of `-max-
autosize` is automatically reset to the supported maximum without
returning an error. The value for `-max-autosize` supports relative
sizing.

``[-min-autosize {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Autosize::
This parameter specifies the minimum size to which the volume can
automatically shrink. If the volume was created with the ``_grow_shrink_``
autosize mode enabled, then the default minimum size is equal to the
initial volume size. If the value of the `-min-autosize` parameter is
invalidated by a manual volume resize, the minimum size is reset to the
volume size. The value for `-min-autosize` supports relative sizing.

``[-autosize-grow-threshold-percent <percent>]`` - Autosize Grow Threshold
Percentage::
This parameter specifies the used space threshold for the automatic growth
of the volume. When the volume's used space becomes greater than this
threshold, the volume will automatically grow unless it has reached the
```

```
maximum autosize.

``[-autosize-shrink-threshold-percent <percent>]`` - Autosize Shrink
Threshold Percentage::
This parameter specifies the used space threshold for the automatic
shrinking of the volume. When the amount of used space in the volume drops
below this threshold, the volume will shrink unless it has reached the
specified minimum size.

``[-autosize-mode {off|grow|grow_shrink}]`` - Autosize Mode::
This parameter specifies the autosize mode for the volume. The supported
autosize modes are:
+
  * ``_off_`` - The volume will not grow or shrink in size in response to
the amount of used space.
  * ``_grow_`` - The volume will automatically grow when used space in the
volume is above the grow threshold.
  * ``_grow_shrink_`` - The volume will grow or shrink in size in response
to the amount of used space.

+
By default, `-autosize-mode` is ``_off_`` for new volumes, except for DP
mirrors, for which the default value is ``_grow_shrink_`` . The ``_grow_``
and ``_grow_shrink_`` modes work together with snapshot autodelete to
automatically reclaim space when a volume is about to become full. The
volume parameter `-space-mgmt-try-first` controls the order in which these
two space reclamation policies are attempted.

| ``[-autosize-reset <true>]`` - Autosize Reset }::
This allows the user to reset the values of autosize, max-autosize, min-
autosize, autosize-grow-threshold-percent, autosize-shrink-threshold-
percent and autosize-mode to their default values. For example, the max-
autosize value will be set to 120% of the current size of the volume.

``[-files <integer>]`` - Total User Visible Files::
This optionally specifies the total number of files for user-visible data
permitted on the volume. This value can be raised or lowered. Raising the
total number of files does not immediately cause additional disk space to
be used to track files. Instead, as more files are created on the volume,
the system dynamically increases the number of disk blocks that are used
to track files. The space assigned to track files is never freed, and the
`files` value cannot be decreased below the current number of files that
can be tracked within the assigned space for the volume.

``[-files-set-maximum {true|false}]`` - Set Total Files (for user-visible
data) to the Highest Value that the Volume can Hold [.advaced-
```

```
priv]#(privilege: advanced)#::
This optionally specifies whether the volume's total number of files will
be set to the highest possible value. If ``_true_`` , the volume's total
number of files is set to the highest value that the volume can hold. Only
``_true>_`` is a valid input. ``_false>_`` is not permitted. To modify the
total number of files to a specific value, use option files.

``[-maxdir-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Directory Size
[.advaced-priv]#(privilege: advanced)#::
This optionally specifies the maximum directory size. The default maximum
directory size is model-dependent, and optimized for the size of system
memory. You can increase it for a specific volume by using this option,
but doing so could impact system performance. If you need to increase the
maximum directory size, work with customer support.

{ ``[-space-slo {none|thick|semi-thick}]`` - Space SLO::
This optionally specifies the Service Level Objective for space management
(the space SLO setting) for the volume. The space SLO value is used to
enforce volume settings so that sufficient space is set aside to meet the
space SLO. The default setting is `none` . There are three supported
values: ``_none_`` , ``_thick_`` and ``_semi-thick_`` .
+
  * `none` : The value of ``_none_`` does not provide any guarantee for
overwrites or enforce any restrictions. It should be used if the admin
plans to manually manage space consumption in the volume and aggregate,
and out of space errors.
  * `thick` : The value of ``_thick_`` guarantees that the hole fills and
overwrites to space-reserved files in this volume will always succeed by
reserving space. To meet this space SLO, the following volume-level
settings are automatically set and cannot be modified:
+
    * Space Guarantee: ``_volume_`` - The entire size of the volume is
preallocated in the aggregate. Changing the volume's `space-guarantee`
type is not supported.
    * Fractional Reserve: ``_100_`` - 100% of the space required for
overwrites is reserved. Changing the volume's `fractional-reserve` setting
is not supported.


  * `semi-thick` : The value of ``_semi-thick_`` is a best-effort attempt
to ensure that overwrites succeed by restricting the use of features that
share blocks and auto-deleting backups and snapshots in the volume. To
meet this space SLO, the following volume-level settings are automatically
set and cannot be modified:
+
    * Space Guarantee: ``_volume_`` - The entire size of the volume is
```

```
preallocated in the aggregate. Changing the volume's `space-guarantee`
type is not supported.
    * Fractional Reserve: ``_0_`` - No space will be reserved for
overwrites by default. However, changing the volume's `fractional-reserve`
setting is supported. Changing the setting to 100 means that 100% of the
space required for overwrites is reserved.
    * Snapshot Autodelete: ``_enabled_`` - Automatic deletion of snapshots
is enabled to reclaim space. To ensure that the overwrites can be
accommodated when the volume reaches threshold capacity, the following
volume snapshot autodelete parameters are set automatically to the
specified values and cannot be modified:
+
      * `enabled` : ``_true_``
      * `commitment` : ``_destroy_``
      * `trigger` : ``_volume_``
      * `defer-delete` : ``_none_``
      * `destroy-list` : ``_vol_clone_`` , ``_lun_clone_`` ,
``_file_clone_`` , ``_cifs_share_``




+
 In addition, with a value of ``_semi-thick_`` , the following
technologies are not supported for the volume:
+
      * File Clones with autodelete disabled: Only full file clones of
files, LUNs or NVMe namespaces that can be autodeleted can be created in
the volume. The use of `autodelete` for file clone create is required.
      * Partial File Clones: Only full file clones of files or LUNs that
can be autodeleted can be created in the volume. The use of `range` for
file clone create is not supported.
      * Volume Efficiency: Enabling volume efficiency is not supported to
allow autodeletion of snapshots.




| ``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
This option controls whether the volume is guaranteed some amount of space
in the aggregate. The default setting for the volumes on All Flash FAS
systems is ``_none_`` , otherwise the default setting is ``_volume_`` .
The `file` setting is no longer supported. Volume guaranteed means that
the entire size of the volume is preallocated. The none value means that
no space is preallocated, even if the volume contains space-reserved files
or LUNs; if the aggregate is full, space is not available even for space-
```

```
reserved files and LUNs within the volume. Setting this parameter to none
enables you to provision more storage than is physically present in the
aggregate (thin provisioning). When you use thin provisioning for a
volume, it can run out of space even if it has not yet consumed its
nominal size and you should carefully monitor space utilization to avoid
unexpected errors due to the volume running out of space. For flexible
root volumes, to ensure that system files, log files, and cores can be
saved, the space-guarantee must be volume. This is to ensure support of
the appliance by customer support, if a problem occurs. Disk space is
preallocated when the volume is brought online and, if not used, returned
to the aggregate when the volume is brought offline. It is possible to
bring a volume online even when the aggregate has insufficient free space
to preallocate to the volume. In this case, no space is preallocated, just
as if the none option had been selected. In this situation, the vol
options and vol status command display the actual value of the space-
guarantee option, but indicate that it is disabled.

``[-fractional-reserve <percent>]`` - Fractional Reserve }::
This option changes the amount of space reserved for overwrites of
reserved objects (LUNs, files) in a volume. The option is set to 100 by
default with `guarantee` set to ``_volume_`` . A setting of 100 means that
100% of the required reserved space is actually reserved so the objects
are fully protected for overwrites. The value is set to 0 by default with
`guarantee` set to ``_none_`` . The value can be either 0 or 100 when
`guarantee` is set to ``_volume_`` or ``_none_`` . Using a value of 0
indicates that no space will be reserved for overwrites. This returns the
extra space to the available space for the volume, decreasing the total
amount of space used. However, this does leave the protected objects in
the volume vulnerable to out of space errors. If the percentage is set to
0%, the administrator must monitor the space usage on the volume and take
corrective action.

``[-min-readahead {true|false}]`` - Minimum Read Ahead [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether minimum readahead is used on the volume.
The default setting is `false` .

``[-atime-update {true|false}]`` - Access Time Update Enabled [.advaced-
priv]#(privilege: advanced)#::
This optionally specifies whether the access time on inodes is updated
when a file is read. The default setting is `true` .

``[-snapdir-access {true|false}]`` - Snapshot Directory Access Enabled::
This optionally specifies whether clients have access to .snapshot
directories. The default setting is `true` .
```

```
``[-percent-snapshot-space <percent>]`` - Space Reserved for Snapshot
Copies::
This optionally specifies the amount of space that is reserved on the
volume for snapshots. The default setting is 5 percent.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
This optionally specifies the snapshot policy for the volume. The default
is the snapshot policy for all volumes on the SVM, as specified by the `-
snapshot-policy` parameter of the xref:{relative_path}vserver-
create.html[vserver create] and xref:{relative_path}vserver-
modify.html[vserver modify] commands. When replacing a snapshot policy on
a volume, any existing snapshots on the volume that do not match any of
the prefixes of the new snapshot policy will not be deleted. This is
because the snapshot scheduler will not clean up older snapshots if the
prefixes do not match. After the new snapshot policy takes effect,
depending on the new retention count, any existing snapshots that continue
to use the same prefixes might be deleted. For example, your existing
snapshot policy is set up to retain 150 weekly snapshots and you create a
new snapshot policy that uses the same prefixes but changes the retention
count to 50 snapshots. After the new snapshot policy takes effect, it will
start deleting older snapshots until there are only 50 remaining.

``[-language <Language code>]`` - Language::
 Use this parameter to change the volume language from *.UTF-8 to utf8mb4.
To change the language of a volume, contact technical support.

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes. This
parameter applies only to FlexGroup volumes. For FlexVol volumes, the
command always runs in the foreground.

``[-nvfail {on|off}]`` - NVFAIL Option::
 Setting this optional parameter to true causes the volume to set the in-
nvfailed-state flag to true, if committed writes to the volume are lost
due to a failure. The in-nvfailed-state flag fences the volume from
further data access and prevents possible corruption of the application
data. Without specifying a value, this parameter is automatically set to
false.

``[-in-nvfailed-state {true|false}]`` - Volume's NVFAIL State [.advaced-
priv]#(privilege: advanced)#::
 This field is automatically set to true on a volume when committed writes
to the volume are possibly lost due to a failure, and the volume has the
nvfail option enabled. With this field set, the client access to the
```

```
volume is fenced to protect against possible corruptions that result from
accessing stale data. The administrator needs to take appropriate recovery
actions to recover the volume from the possible data loss. After the
recovery is completed, the administrator can clear this field and restore
the client access to the volume. This field can be cleared using the CLI
but it cannot be set.

``[-dr-force-nvfail {on|off}]`` - Force NVFAIL on MetroCluster
Switchover::
 Setting this optional parameter to true on a volume causes the
MetroCluster switchover operation to set the in-nvfailed-state flag to
true on that volume. The in-nvfailed-state flag prevents further data
access to the volume. The default value is false. This parameter has no
effect on a negotiated or an automatic switchover.

``[-filesys-size-fixed {true|false}]`` - Is File System Size Fixed::
This parameter is only applicable for a DP relationship. This option
causes the file system to remain the same size and not grow or shrink when
a SnapMirrored volume relationship is broken, or when a volume add is
performed on it. It is automatically set to true when a volume becomes a
SnapMirrored volume. It stays set to true after the SnapMirror break
command is issued for the volume. This allows a volume to be SnapMirrored
back to the source without needing to add disks to the source volume. If
the volume is a flexible volume and the volume size is larger than the
file system size, setting this option to false forces the volume size to
equal the file system size. The default setting is false.

``[-extent-enabled {off|on|space-optimized}]`` - (DEPRECATED)-Extent
Option::
+
NOTE: This parameter has been deprecated and may be removed in a future
release of ONTAP.
+
 Setting this option to `on` or `space-optimized` enables extents in the
volume. This causes application writes to be written in the volume as a
write of a larger group of related data blocks called an extent. Using
extents may help workloads that perform many small random writes followed
by large sequential reads. However, using extents may increase the amount
of disk operations performed on the controller, so this option should only
be used where this trade-off is desired. If the option is set to `space-
optimized` then the reallocation update will not duplicate blocks from
snapshots into the active file system, and will result in conservative
space utilization. Using `space-optimized` may be useful when the volume
has snapshots or is a SnapMirror source, when it can reduce the storage
used in the volume and the amount of data that SnapMirror needs to move on
the next update. The `space-optimized` value can result in degraded read
```

```
performance of snapshots. The default value is `off` ; extents are not
used.

``[-space-mgmt-try-first {volume_grow|snap_delete}]`` - Primary Space
Management Strategy::
A flexible volume can be configured to automatically reclaim space in case
the volume is about to run out of space, by either increasing the size of
the volume using autogrow or deleting snapshots in the volume using
snapshot autodelete. If this option is set to volume_grow the system will
try to first increase the size of volume before deleting snapshots to
reclaim space. If the option is set to snap_delete the system will first
automatically delete snapshots and in case of failure to reclaim space
will try to grow the volume.

``[-read-realloc {off|on|space-optimized}]`` - Read Reallocation Option::
Setting this option to `on` or `space-optimized` enables read reallocation
in the volume. This results in the optimization of file layout by writing
some blocks to a new location on disk. The layout is updated only after
the blocks have been read because of a user read operation, and only when
updating their layout will provide better read performance in the future.
Using read reallocation may help workloads that perform a mixture of
random writes and large sequential reads. If the option is set to `space-
optimized` then the reallocation update will not duplicate blocks from
snapshots into the active file system, and will result in conservative
space utilization. Using `space-optimized` may be useful when the volume
has snapshots or is a SnapMirror source, when it can reduce the storage
used in the volume and the amount of data that SnapMirror needs to move on
the next update. The `space-optimized` value can result in degraded read
performance of snapshots. The default value is `off` .

``[-sched-snap-name {create-time|ordinal}]`` - Naming Scheme for Automatic
Snapshot Copies::
This option specifies the naming convention for automatic snapshots. If
set to ``_create-time_`` , automatic snapshots are named using the format
``_<schedule_name>.yyyy-mm-dd_hhmm_`` . Example: "hourly.2010-04-01_0831".
If set to ``_ordinal_`` , only the latest automatic snapshot is named
using the format ``_<schedule_name>.<n>_`` . Example: "hourly.0". Older
automatic snapshots are named using the format ``_<schedule_name>.yyyy-mm-
dd_hhmm_`` . Example: "hourly.2010-04-01_0831".

{ ``[-qos-policy-group <text>]`` - QoS Policy Group Name::
This optional parameter specifies which QoS policy group to apply to the
volume. This policy group defines measurable service level objectives
(SLOs) that apply to the storage objects with which the policy group is
associated. If you do not assign a policy group to a volume, the system
will not monitor and control the traffic to it. To remove this volume from
```

```
a policy group, enter the reserved keyword "none".

| ``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name
}::
This optional parameter specifies which QoS adaptive policy group to apply
to the volume. This policy group defines measurable service level
objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on
the volume allocated space or used space. To remove this volume from an
adaptive policy group, enter the reserved keyword "none".

``[-caching-policy <text>]`` - Caching Policy Name::
This parameter specifies the caching policy to apply to the volume. A
caching policy defines how the system caches this volume's data in a Flash
Pool aggregate or Flash Cache modules.
+
Both metadata and user data are eligible for caching. Metadata consists of
directories, indirect blocks and system metafiles. They are eligible for
read caching only. When a random write pattern is detected on user data,
the first such write is eligible for read caching while all subsequent
overwrites are eligible for write caching. The available caching policies
are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.
  * noread-random_write - Write caches all randomly overwritten user data
blocks. It does not do any read caching.
  * meta-random_write - Read caches all metadata and write caches randomly
overwritten user data blocks.
  * random_read_write-random_write - Read caches all metadata, randomly
read and randomly written user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read-random_write - Read caches all metadata, randomly read and
sequentially read user data blocks. It also write caches randomly
overwritten user data blocks.
```

```
  * all_read_random_write-random_write - Read caches all metadata,
randomly read, sequentially read and randomly written user data. It also
write caches randomly overwritten user data blocks.
  * all-random_write - Read caches all data blocks read and written. It
also write caches randomly overwritten user data blocks.

+
Note that in a caching-policy name, a hyphen (-) separates read and write
policies. Default caching-policy is auto.

``[-is-autobalance-eligible {true|false}]`` - Is Eligible for Auto Balance
Aggregate [.advaced-priv]#(privilege: advanced)#::
If the Auto Balance feature is enabled, this parameter specifies whether
the volume might be considered for system workload balancing. When set to
``_true_`` , the Auto Balance Aggregate feature might recommend moving
this volume to another aggregate. The default value is ``_true_`` .

``[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum size of
a FlexGroup Constituent [.advaced-priv]#(privilege: advanced)#::
This optionally specifies the maximum size of a FlexGroup constituent. The
default value is determined by checking the maximum FlexVol size setting
on all nodes used by the FlexGroup. The smallest value found is selected
as the default for the `-max-constituent-size` for the FlexGroup. This
parameter applies to FlexGroup volumes only.

``[-vserver-dr-protection {protected|unprotected}]`` - Vserver DR
Protection::
 This optionally specifies whether the volume should be protected by
Vserver level SnapMirror. This parameter is applicable only if the Vserver
is the source of a Vserver level SnapMirror relationship.

``[-is-space-reporting-logical {true|false}]`` - Logical Space Reporting::
This optionally specifies whether to report space logically on the volume.
When space is reported logically, ONTAP reports the volume space such that
all the physical space saved by the storage efficiency features are also
as reported as used. The default setting is ``_false_`` .

``[-is-space-enforcement-logical {true|false}]`` - Logical Space
Enforcement::
This optionally specifies whether to perform logical space accounting on
the volume. When space is enforced logically, ONTAP enforces volume
settings such that all the physical space saved by the storage efficiency
features will be calculated as used. The default setting is ``_false_`` .

``[-tiering-policy <Tiering Policy>]`` - Volume Tiering Policy::
This optional parameter specifies the tiering policy to apply to the
```

```
volume. This policy determines whether the user data blocks of a volume in
a FabricPool will be tiered to the cloud tier when they become cold.
FabricPool combines Flash (performance tier) with an object store (cloud
tier) into a single aggregate. The temperature of a volume block increases
if it is accessed frequently, and it decreases when it is not.
+
The available tiering policies are:
+
  * snapshot-only - This policy allows tiering of only the volume
snapshots not associated with the active file system. The default cooling
period is 2 days. The `-tiering-minimum-cooling-days` parameter can be
used to override the default.
  * auto - This policy allows tiering of both snapshot data and active
file system user data to the cloud tier. The default cooling period is 31
days. The `-tiering-minimum-cooling-days` parameter can be used to
override the default.
  * none - Volume blocks will not be tiered to the cloud tier.
  * all - This policy allows tiering of both snapshot data and active file
system user data to the cloud tier as soon as possible without waiting for
a cooling period. On DP volumes, this policy allows all transferred user
data blocks to start in the cloud tier.



``[-cloud-retrieval-policy {default|on-read|never|promote}]`` - Volume
Cloud Retrieval Policy [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the cloud retrieval policy for the
volume. This policy determines which tiered out blocks to retrieve from
the capacity tier to the performance tier.
+
The available cloud retrieval policies are:
+
  * default - This policy retrieves tiered data based on the underlying
tiering policy. If the tiering policy is 'auto', tiered data is retrieved
only for random client driven data reads. If the tiering policy is 'none'
or 'snapshot-only', tiered data is retrieved for random and sequential
client driven data reads. If the tiering policy is 'all', tiered data is
not retrieved.
  * on-read - This policy retrieves tiered data for all client driven data
reads.
  * never - This policy never retrieves tiered data.
  * promote - This policy retrieves all eligible tiered data automatically
during the next scheduled scan. It is only supported when the tiering
policy is 'none' or 'snapshot-only'. If the tiering policy is 'snapshot-
only', the only data brought back is the data in the AFS. Data that is
only in a snapshot stays in the cloud.
```

```
``[-tiering-minimum-cooling-days <integer>]`` - Volume Tiering Minimum
Cooling Days [.advaced-priv]#(privilege: advanced)#::
 This parameter specifies the minimum number of days that user data blocks
of the volume must be cooled before they can be considered cold and tiered
out to the cloud tier. For volumes hosted on FabricPools, this parameter
is used for tiering purposes and does not affect the reporting of inactive
data. For volumes hosted on non-FabricPools, this parameter affects the
cooling window used for reporting inactive data. The value specified
should be greater than the frequency with which applications in the volume
shift between different sets of data. Valid values are between 2 and 183.
This parameter cannot be set when volume tiering policy is either "none"
or "all".

``[-tiering-object-tags <text>,...]`` - Tags to be Associated with Objects
Stored on a FabricPool::
This optional parameter specifies tiering object tags to be associated
with objects stored on a FabricPool.
+
Object tags should follow these rules:
+
  * Each object tag should be a key-value pair separated by '='.
  * Multiple tags should be separated by ','. Overall tags should be in
format key1=value1[,key2=value2,...].
  * All tags of a volume must have a unique key.
  * Each tag key should start with either a letter or an underscore. Keys
should contain only alphanumeric characters and underscores. Maximum
allowed limit is of 127 characters.
  * Each tag value should be of maximum 127 characters consisting of only
alphanumeric characters and underscores.
  * Maximum 4 object tags are allowed per volume.
  * To remove existing tags of a volume, specify empty list as a parameter
value.



``[-anti-ransomware-state {disabled|enabled|dry-run|paused|dry-run-
paused|enable-paused|disable-in-progress}]`` - Anti-ransomware State::
Use this parameter to modify the anti-ransomware-state of the volume. The
states dry-run-paused, enable-paused and disable-in-progress are non-
modifiable. The following states are supported for modify:
+
  * ``_enabled_`` - the feature is enabled.
  * ``_disabled_`` - the feature is disabled.
```

```
  * ``_dry-run_`` - the feature is in evaluation mode.
  * ``_paused_`` - the feature is in pause mode.



``[-granular-data {disabled|basic|advanced}]`` - Granular data::
This parameter specifies the type of granular data storage supported on
the volume. Once this property is changed to any value other than
``_disabled_`` on a volume, it cannot be disabled on that volume using
this parameter. It can only be disabled on that volume by restoring a
snapshot. This property is only supported on FlexGroup volumes. Files
modified by the granular data algorithms use granular data, which might
impact read/write I/O performance.
+
The available options are:
+
  * disabled - No volume capacity rebalancing operations are performed.
  * basic - Volume capacity rebalancing operations are enabled.
  * advanced - All ``_basic_`` features are enabled. In addition, granular
data distribution is enabled. Granular data distribution allows for
improved capacity rebalancing and usage balancing in certain scenarios.
This property can be set when the current value for the volume is
``_disabled_`` or ``_basic_`` . Once this property is set on a volume, it
cannot be modified to ``_disabled_`` or ``_basic_`` using this parameter.
It can only be changed on that volume by restoring a snapshot.



``[-atime-update-period <integer>]`` - Access Time Update Period (Seconds)
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the time period that must be elapsed in between
consecutive atime update events. Note: The value of this parameter is only
used if `-atime-update` is enabled on the volume.

``[-snapshot-locking-enabled {true|false}]`` - Enable Snapshot Copy
Locking::
This parameter specifies whether snapshot locking is enabled, based on an
independent compliance-clock time. Once this property is enabled, it
cannot be disabled and the volume cannot be be deleted until all locked
snapshots are past their expiry time.

``[-is-large-size-enabled {true|false}]`` - Are Large Size Volumes and
Files Enabled::
This parameter specifies whether support for large Flexible volumes and
large files is enabled. If this value is true, the maximum volume size for
a Flexible volume will be 300TB and the maximum size of a single file will
```

```
be 128TB. If this value is false, the maximum volume size for a Flexible
volume will be 100TB and the maximum size of a single file will be 16TB.
The default value is false.

``[-is-preserve-unlink-enabled {true|false}]`` - Is Preserve Unlink
Enabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether preserve unlink is enabled for NFSV4.1
shares on the volume. Once this feature is enabled, a file removal op on a
file with existing share locks will result in the file being moved to the
trash directory. Clients will continue to be able to access the file using
its file handle. The last close for that file will result in deletion of
the file. The default value is false.

``[-is-cloud-write-enabled {true|false}]`` - Is Cloud Write Enabled
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether cloud write is enabled on the volume.
This feature is only available on volumes in FabricPools. The default
value is false.

``[-aggressive-readahead-mode
{none|file_prefetch|sequential_read|cross_file_sequential_read}]`` -
Aggressive readahead mode [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the aggressive readahead mode of the volume. When
set to file_prefetch, on a file read, the system aggressively issues
readaheads for all of the blocks in the file and retains those blocks in a
cache for a finite period of time. This feature is only available on
FabricPool volumes on FSx for ONTAP and Cloud Volumes ONTAP. The default
value is none.

== Examples

The following example modifies a volume named vol4 on a Vserver named vs0.
The volume's export policy is changed to default_expolicy and its size is
changed to 500 GB.
```

cluster1::> volume modify -vserver vs0 -volume vol4 -policy default_expolicy -size 500g

```
The following example modifies a volume named vol2. It enables autogrow
and sets the maximum autosize to 500g
```

cluster1::> volume modify -volume vol2 -autosize-mode grow -max-autosize 500g

```
The following example modifies a volume named vol2 to have a space
guarantee of none.
```

cluster1::> volume modify -space-guarantee none -volume vol2

```
The following example modifies all volumes in Vserver vs0 to have a
fractional reserve of 30%.
```

cluster1::> volume modify -fractional-reserve 30 -vserver vs0 *

```
The following example modifies a volume named vol2 to grow in size by 5
gigabytes
```

cluster1::> volume modify -volume vol2 -size +5g

```
The following example modifies a volume named vol2 to have a different
caching policy. The volume must be on a Flash Pool aggregate.
```

cluster1::> volume modify -volume vol2 -caching-policy none

```
== Related Links

* xref:{relative_path}volume-rename.html[volume rename]
* xref:{relative_path}volume-make-vsroot.html[volume make-vsroot]
* xref:{relative_path}vserver-export-policy-create.html[vserver export-
policy create]
* xref:{relative_path}vserver-create.html[vserver create]
* xref:{relative_path}vserver-modify.html[vserver modify]




[[ID0f3c76a4c6432cc681d95f8228cf9bf4]]
= volume mount
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Mount a volume on another volume with a junction-path

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `volume mount` command mounts a volume at a specified junction path.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the volume that is to be mounted.

``-junction-path <junction path>`` - Junction Path Of The Mounting
Volume::
This specifies the junction path of the mounted volume. The junction path
name is case insensitive and must be unique within a Vserver's namespace.

``[-active {true|false}]`` - Activate Junction Path::
This optionally specifies whether the mounted volume is accessible. The
default setting is `false` . If the mounted path is not accessible, it
does not appear in the Vserver's namespace.

``[-policy-override {true|false}]`` - Override The Export Policy::
This optionally specifies whether the parent volume's export policy
overrides the mounted volume's export policy. The default setting is
`false` .

== Examples

The following example mounts a volume named user_tsmith on a Vserver named
vs0. The junction path for the mounted volume is /user/tsmith. The mounted
volume is accessible, and the mounted volume's export policy is not
overridden by the parent volume's export policy.
```

node::> volume mount -vserver vs0 -volume user_tsmith
-junction-path /user/tsmith -active true -policy-override false

```
[[ID973756d63fefa06727d7e52adff0ee5f]]
= volume offline
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Take an existing volume offline

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

The `volume offline` command takes the volume offline. If the volume is
already in restricted or iron_restricted state, then it is already
unavailable for data access, and much of the following description does
not apply. The current root volume might not be taken offline. Several
operations being performed on the volume in question can prevent the
volume offline command from succeeding for various lengths of time. If
such operations are required, the command might take additional time to
complete. If you set the volume junction path, it is unmounted when the
volume is taken offline. The junction path is not restored when the volume
is brought online again. When the volume offline command fails, the
command is aborted. The `-force` flag can be used to forcibly offline a
volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver from which the volume is to be
taken offline. If only one data Vserver exists, you do not need to specify
this parameter.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume that is to be taken offline.

``[-f, -force <true>]`` - Force Offline::
This specifies whether the offline operation is forced. Using this option
to force a volume offline can potentially disrupt access to other volumes.
The default setting is `false` .

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes. This
parameter applies only to FlexGroup volumes. For FlexVol volumes, the
command always runs in the foreground.

``[-disable-block-storage-check <true>]`` - Disable Check for Existing
LUNs and NVMe Namespaces::
Taking the volume offline will make all associated LUNs and NVMe over
Fabrics (NVMe-oF) namespaces unavailable, which normally requires a user
confirmation. If this parameter is specified, the command proceeds without
a confirmation. The default setting is `false`
```

```
== Examples

The following example takes the volume named vol1 offline:
```

cluster1::> volume offline vol1
Volume 'vs1:vol1' is now offline.

```
[[ID2e419998e9e52f567605c11222cb2de0]]
= volume online
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Bring an existing volume online

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume online` command brings the volume online. A volume can only be
brought online if it is offline or restricted. If the volume is
inconsistent but has not lost data, the user will be cautioned and
prompted before bringing it online. It is advisable to run wafl-iron (or
do a xref:{relative_path}snapmirror-initialize.html[snapmirror initialize]
in case of a replica volume) prior to bringing an inconsistent volume
online. Bringing an inconsistent volume online increases the risk of
further file system corruption. If the containing aggregate cannot honor
the space guarantees required by this volume, the volume online operation
will fail. It is not advisable to use volumes with their space guarantees
disabled. Lack of free space can lead to failure of writes which in turn
can appear as data loss to some applications.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the name of the Vserver from which the volume is
to be brought online. If only one data Vserver exists, you do not need to
specify this parameter.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume that is to be brought
online.
```

```
``[-f, -force <true>]`` - Force Online::
When this parameter is used, the volume will be brought online even if
there is not enough space available in the aggregate to honor the volume's
space guarantee.

``[-foreground {true|false}]`` - Foreground Process::
This parameter specifies whether the operation runs in the foreground. The
default setting is `true` (the operation runs in the foreground). When set
to true, the command will not return until the operation completes. This
parameter applies only to FlexGroup volumes. For FlexVol volumes, the
command always runs in the foreground.

== Examples

The following example brings a volume named vol1 online:
```

cluster1::> volume online vol1
Volume 'vs1:vol1' is now online.

```
== Related Links

* xref:{relative_path}snapmirror-initialize.html[snapmirror initialize]




[[ID1166c9664f0813e218f7104bfe1ddd35]]
= volume rehost
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rehost a volume from one Vserver into another Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `volume rehost` command rehosts a volume from source Vserver onto
destination Vserver. The volume name must be unique among the other
volumes on the destination Vserver. This command is not supported in a
MetroCluster configuration.
```

```
== Parameters

``-vserver <vserver name>`` - Source Vserver name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Target volume name::
This specifies the volume that is to be rehosted.

``-destination-vserver <vserver name>`` - Destination Vserver name::
This specifies the destination Vserver where the volume must be located
post rehost operation.

{ ``[-force-unmap-luns {true|false}]`` - Unmap LUNs in volume::
This specifies whether the rehost operation should unmap LUNs present on
volume. The default setting is ``_false_`` (the rehost operation shall not
unmap LUNs). When set to true, the command will unmap all mapped LUNs on
the volume.

| ``[-auto-remap-luns {true|false}]`` - Automatic Remap of LUNs }::
This specifies whether the rehost operation should perform LUN mapping
operation at the destination Vserver for the LUNs mapped on the volume at
the source Vserver. The default setting is ``_false_`` (the rehost
operation shall not map LUNs at the destination Vserver). When set to
true, at the destination Vserver the command will create initiators groups
along with the initiators (if present) with same name as that of source
Vserver. Then the LUNs on the volume are mapped to initiator groups at the
destination Vserver as mapped in source Vserver.

``[-key-manager-attribute <text>]`` - Optional CRN::
This optional parameter specifies an additional key manager attribute that
should be an identifier-value pair separated by '=', ex. CRN=my-unique-
value
+
The following identifiers are currently supported:
+
  * CRN



== Examples

The following example rehosts a volume named vol3 on Vserver named vs1 to
a destination Vserver named vs2:
```

cluster::> volume rehost -vserver vs1 -volume vol3 -destination-vserver vs2

```
[[ID1d962368adae272762573528febdaba7]]
= volume rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename an existing volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume rename` command renames a volume. The volume name must be
unique among the other volumes in the same Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located. For a node's
root volume, use the name of the node for this parameter.

``-volume <volume name>`` - Volume Name::
This specifies the volume that is to be renamed.

``-newname <volume name>`` - Volume New Name::
This specifies the volume's new name. A volume's name must start with an
alphabetic character (a to z or A to Z) and be 203 or fewer characters in
length.

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes. This
parameter applies only to FlexGroup volumes. For FlexVol volumes, the
command always runs in the foreground.

== Examples

The following example renames a volume named vol3_backup as vol3_save on a
Vserver named vs0:
```

node::> volume rename -vserver vs0 -volume vol3_backup -newname vol3_save

```
[[IDb7be188821b58ae9ba7e866d6668d70f]]
= volume restrict
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restrict an existing volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume restrict` command puts the volume in restricted state. If the
volume is online, then it will be made unavailable for data access as
described under xref:{relative_path}volume-offline.html[volume offline] .

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver from which the volume is to be
restricted. If only one data Vserver exists, you do not need to specify
this parameter.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume that is to be restricted.

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes. This
parameter applies only to FlexGroup volumes. For FlexVol volumes, the
command always runs in the foreground.

== Examples

The following example restricts a volume named vol1:
```

cluster1::> volume restrict vol1
Volume 'vs1:vol1' is now restricted.

```
== Related Links
```

```
* xref:{relative_path}volume-offline.html[volume offline]




[[ID835d2a96dcb6579cc3c8ea491ed8602e]]
= volume show-footprint
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of volumes and their data and metadata footprints in their
associated aggregate.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume show-footprint` command displays information about the space
used in associated aggregates by volumes and features enabled in volumes.
The command output depends on the parameter or parameters specified with
the command. If no parameters are specified, the command displays the
following information about all volumes.  + + If the associated aggregates
have an object-store attached to them, then the command displays a more
detailed split up of the space used in each tier. This additional
information is useful to show per-tier space usage which can be used to
estimate the space requirements and transfer duration when moving a volume
to a different tier with `volume move` .

== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter and the `-volume` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
```

```
volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter and the `-vserver` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
all volumes matching the specified name.

``[-volume-msid <integer>]`` - Volume MSID::
If this parameter is specified, the command displays information only
about the volume that has the specified MSID.

``[-volume-dsid <integer>]`` - Volume DSID::
If this parameter is specified, the command displays information only
about the volume that has the specified DSID.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
If this parameter is specified, the command displays information only
about the volume on the vserver which has the specified UUID.

``[-aggregate <aggregate name>]`` - Aggregate Name::
If this parameter is specified, the command displays information only
about the volumes that are associated with the specified aggregate.

``[-aggregate-uuid <UUID>]`` - Aggregate UUID::
If this parameter is specified, the command displays information only
about the volumes on the aggregate which have the specified UUID.

``[-hostname <text>]`` - Hostname::
If this parameter is specified, the command displays information only
about the volumes that belong to the specified host.

``[-tape-backup-metafiles-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Tape
Backup Metadata Footprint::
If this parameter is specified, the command displays information only
about the volumes whose tape backup metafiles use the specified amount of
space in the aggregate.

``[-tape-backup-metafiles-footprint-percent <percent>]`` - Tape Backup
Metadata Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose tape backup metafiles use the specified percentage
of space in the aggregate.

``[-dedupe-metafiles-footprint {<integer>[KB|MB|GB|TB|PB]}]`` -
Deduplication Footprint::
If this parameter is specified, the command displays information only
```

```
about the volumes whose deduplication metafiles use the specified amount
of space in the aggregate.

``[-dedupe-metafiles-footprint-percent <percent>]`` - Deduplication
Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose deduplication metafiles use the specified
percentage of space in the aggregate.

``[-dedupe-metafiles-temporary-footprint {<integer>[KB|MB|GB|TB|PB]}]`` -
Temporary Deduplication Footprint::
If this parameter is specified, the command displays information only
about the volumes whose temporary deduplication metafiles use the
specified amount of space in the aggregate.

``[-dedupe-metafiles-temporary-footprint-percent <percent>]`` - Temporary
Deduplication Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose temporary deduplication metafiles use the
specified percentage of space in the aggregate.

``[-cross-volume-dedupe-metafiles-footprint {<integer>[KB|MB|GB|TB|PB]}]``
- Cross Volume Deduplication Footprint::
If this parameter is specified, the command displays information only
about the volumes whose cross volume deduplication metafiles use the
specified amount of space in the aggregate.

``[-cross-volume-dedupe-metafiles-footprint-percent <percent>]`` - Cross
Volume Deduplication Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose cross volume deduplication metafiles use the
specified percentage of space in the aggregate.

``[-cross-volume-dedupe-metafiles-temporary-footprint
{<integer>[KB|MB|GB|TB|PB]}]`` - Cross Volume Temporary Deduplication
Footprint::
If this parameter is specified, the command displays information only
about the volumes whose cross volume deduplication temporary metafiles use
the specified amount of space in the aggregate.

``[-cross-volume-dedupe-metafiles-temporary-footprint-percent <percent>]``
- Cross Volume Temporary Deduplication Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose cross volume deduplication temporary metafiles use
the specified percentage of space in the aggregate.
```

```
``[-volume-blocks-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Data
Footprint::
If this parameter is specified, the command displays information only
about the volumes whose data blocks use the specified amount of space in
the aggregate.
+
This field is the total amount of data written to the volume. It includes
data in the active file system in the volume as well as data that is
consumed by volume snapshots. This row only includes data and not reserved
space, so when volumes have reserved files, the volume's total used in the
xref:{relative_path}volume-show-space.html[volume show-space] command
output can exceed the value in this row.

``[-volume-blocks-footprint-percent <percent_no_limit>]`` - Volume Data
Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose data blocks use the specified percentage of space
in the aggregate.

``[-flexvol-metadata-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Flexible
Volume Metadata Footprint::
If this parameter is specified, the command displays information only
about the volumes whose file system metadata uses the specified amount of
space in the aggregate.
+
This field includes the space used or reserved in the aggregate for
metadata associated with this volume.

``[-flexvol-metadata-footprint-percent <percent>]`` - Flexible Volume
Metadata Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose file system metadata uses the specified percentage
of space in the aggregate.

``[-delayed-free-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Delayed Free
Blocks::
If this parameter is specified, the command displays information only
about the volumes whose delayed free blocks use the specified amount of
space in the aggregate.
+
When ONTAP frees space in a volume, this space is not always immediately
shown as free in the aggregate. This is because the operations to free the
space in the aggregate are batched for increased performance. Blocks that
are declared free in the FlexVol volume but which are not yet free in the
aggregate are called "delayed free blocks" until the associated delayed
free blocks are processed. For SnapMirror destinations, this row will have
```

```
a value of 0 and will not be displayed.

``[-delayed-free-footprint-percent <percent>]`` - Delayed Free Blocks
Percent::
If this parameter is specified, the command displays information only
about the volumes that have the specified amount of blocks waiting to be
freed in the aggregate. This space is called "delayed free blocks".

``[-snapmirror-destination-footprint {<integer>[KB|MB|GB|TB|PB]}]`` -
SnapMirror Destination Footprint::
If this parameter is specified, the command displays information only
about the volumes whose SnapMirror transfer uses the specified amount of
space in the aggregate.
+
During a SnapMirror transfer, this row will include incoming SnapMirror
data and SnapMirror-triggered delayed free blocks from previous SnapMirror
transfers.

``[-snapmirror-destination-footprint-percent <percent>]`` - SnapMirror
Destination Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose SnapMirror transfer uses the specified percentage
of space in the aggregate.

``[-volume-guarantee-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Volume
Guarantee::
If this parameter is specified, the command displays information only
about the volumes whose guarantees use the specified amount of space in
the aggregate.
+
This field includes the amount of space reserved by this volume in the
aggregate for future writes. The amount of space reserved depends on the
guarantee type (the provisioning mode) of the volume. + For a "volume"
guaranteed volume, this is the size of the volume minus the amount in the
Volume Data Footprint row. + For a "file" guaranteed volume, this is the
sum of all of the space reserved for hole fills and overwrites in all of
the space reserved files in the volume.

``[-volume-guarantee-footprint-percent <percent>]`` - Volume Guarantee
Percent::
If this parameter is specified, the command displays information only
about the volumes whose guarantees use the specified percentage of space
in the aggregate.

``[-file-operation-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - File
Operation Metadata::
```

```
If this parameter is specified, the command displays information only
about the volumes that have file operation metadata using the specified
amount of space in the aggregate.
+
This metadata is used by file move and copy operations. Although it is not
returned as free space once the operations are complete, it can be reused
by future file operations.

``[-file-operation-metadata-percent <percent>]`` - File Operation Metadata
Percent::
If this parameter is specified, the command displays information only
about the volumes that have file operation metadata using the specified
percentage of space in the aggregate.

``[-total-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Total Footprint::
If this parameter is specified, the command displays information only
about the volumes which use the specified amount of space in the
aggregate. This field is the sum of the other rows in this table.

``[-total-footprint-percent <percent_no_limit>]`` - Total Footprint
Percent::
If this parameter is specified, the command displays information only
about the volumes which use the specified percentage of space in the
aggregate.

``[-aggregate-size {<integer>[KB|MB|GB|TB|PB]}]`` - Containing Aggregate
Size::
If this parameter is specified, the command displays information only
about the volumes that are associated with an aggregate of the specified
size.

``[-bin0-name <text>]`` - Name for bin0::
If this parameter is specified, the command displays information only
about volumes whose associated aggregate has an object store attached to
it with an active file system tier name matches the specified value.

``[-volume-blocks-footprint-bin0 {<integer>[KB|MB|GB|TB|PB]}]`` - Volume
Footprint for bin0::
If this parameter is specified, the command displays information only
about volumes whose space in use in the performance tier of the aggregate
matches the specified value.

``[-volume-blocks-footprint-bin0-percent <percent>]`` - Volume Footprint
bin0 Percent::
If this parameter is specified, the command displays information only
about volumes whose percentage space in use in the performance tier of the
```

```
aggregate matches the specified value.

``[-bin1-name <text>]`` - Name for bin1::
If this parameter is specified, the command displays information only
about volumes whose associated aggregate has a cloud tier attached to it
with a configuration name that matches the specified value.

``[-volume-blocks-footprint-bin1 {<integer>[KB|MB|GB|TB|PB]}]`` - Volume
Footprint for bin1::
If this parameter is specified, the command displays information only
about volumes whose space in use in the first cloud tier bucket attached
to the associated aggregate matches the specified value. This includes the
space used by the blocks in the volume that are staged before being moved
to the cloud tier.

``[-volume-blocks-footprint-bin1-percent <percent>]`` - Volume Footprint
bin1 Percent::
If this parameter is specified, the command displays information only
about volumes whose percentage space in use in the first cloud tier bucket
attached to associated aggregate matches the specified value. This
includes the space used by the blocks in the volume that are staged before
being moved to the cloud tier.

``[-total-dedupe-metafiles-footprint {<integer>[KB|MB|GB|TB|PB]}]`` -
Total Deduplication Footprint::
If this parameter is specified, the command displays information only
about the volumes whose total deduplication metafiles use the specified
amount of space in the aggregate.

``[-total-dedupe-metafiles-footprint-percent <percent>]`` - Total
Deduplication Footprint Percent::
If this parameter is specified, the command displays information only
about the volumes whose total deduplication metafiles use the specified
percentage of space in the aggregate.

``[-auto-adaptive-compression-footprint-data-reduction
{<integer>[KB|MB|GB|TB|PB]}]`` - Footprint Data Reduction by Auto Adaptive
Compression::
If this parameter is specified, the command displays information only
about the volumes whose estimation of footprint data reduction due to auto
adaptive compression use the specified amount of space in the aggregate.

``[-auto-adaptive-compression-footprint-data-reduction-percent
<percent>]`` - Footprint Data Reduction by Auto Adaptive Compression
Percent::
If this parameter is specified, the command displays information only
```

```
about the volumes whose percentage estimation of footprint data reduction
due to auto adaptive compression use the specified amount of space in the
aggregate.

``[-total-footprint-data-reduction {<integer>[KB|MB|GB|TB|PB]}]`` - Total
Footprint Data Reduction::
If this parameter is specified, the command displays information only
about the total footprint data reduction by the volume in the aggregate.

``[-total-footprint-data-reduction-percent <percent>]`` - Total Footprint
Data Reduction Percent::
If this parameter is specified, the command displays information only
about the percentage of total footprint data reduction by the volume in
the aggregate.

``[-capacity-tier-footprint-data-reduction {<integer>[KB|MB|GB|TB|PB]}]``
- Footprint Data Reduction by Capacity Tier::
If this parameter is specified, the command displays information only
about the footprint data reduction in the cloud tier by the volume in a
FabricPool.

``[-capacity-tier-footprint-data-reduction-percent <percent>]`` -
Footprint Data Reduction by Capacity Tier Percent::
If this parameter is specified, the command displays information only
about the percentage of footprint data reduction in the cloud tier bucket
by the volume in a FabricPool.

``[-effective-total-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Effective
Total after Footprint Data Reduction::
If this parameter is specified, the command displays information only
about the total space used by the volume in the aggregate after total
footprint data reduction.

``[-effective-total-footprint-percent <percent>]`` - Effective Total after
Footprint Data Reduction Percent::
If this parameter is specified, the command displays information only
about the percentage of total space used by the volume in the aggregate
after total footprint data reduction.

``[-compaction-footprint-data-reduction {<integer>[KB|MB|GB|TB|PB]}]`` -
Footprint Data Reduction by Compaction::
If this parameter is specified, the command displays information only
about the volumes whose estimation of footprint data reduction due to
compaction use the specified amount of space in the aggregate.

``[-compaction-footprint-data-reduction-percent <percent>]`` - Footprint
```

```
Data Reduction by Compaction Percent::
If this parameter is specified, the command displays information only
about the volumes whose percentage estimation of footprint data reduction
due to compaction use the specified amount of space in the aggregate.

``[-total-metadata-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Total
Metadata Footprint::
If this parameter is specified, the command displays information only
about volume that have the specified amount of space used by volume
metadata in the aggregate. The value of this field will be higher than the
field 'total-metadata' from the xref:{relative_path}volume-show-
space.html[volume show-space] command output since this includes
additional metadata of the volume in the aggregate.

``[-total-metadata-footprint-percent <percent>]`` - Total Metadata
Footprint Percentage::
.If this parameter is specified, the command displays information only
about the volume that have the specified percentage of space used by
volume metadata in the aggregate.

== Examples

The following example displays information about all volumes in the
system.
```

cluster1::> volume show-footprint
Vserver : nodevs
Volume : vol0
Feature Used Used%
--- --- ---
Volume Data Footprint 103.1MB 11%
Volume Guarantee 743.6MB 83%
Flexible Volume Metadata 4.84MB 1%
Delayed Frees 4.82MB 1%
Total Metadata Footprint 10.2MB 2%
Total Footprint 856.3MB 95%
Vserver : thevs
Volume : therootvol
Feature Used Used%
--- --- ---
Volume Data Footprint 116KB 0%
Volume Guarantee 19.83MB 1%
Flexible Volume Metadata 208KB 0%
Delayed Frees 60KB 0%
Total Metadata Footprint 320KB 0%
Total Footprint 20.20MB 1%
Vserver : thevs
Volume : thevol
Feature Used Used%

--- --- ---
Volume Data Footprint 128KB 0%
Volume Guarantee 2.00GB 76%
Flexible Volume Metadata 11.38MB 0%
Delayed Frees 428KB 0%
Total Metadata Footprint 13.52MB 0%
Total Footprint 2.01GB 76%

3 entries were displayed.

```
The following example displays information about all volumes in a system
and highlights a scenario where the aggregates associated with volumes
have a cloud tier attached to them.
```

cluster-1::> vol show-footprint
Vserver : vsim1
Volume : vol0
Feature Used Used%
--- --- ---
Volume Data Footprint 2.57GB 81%
Volume Guarantee 266.1MB 8%
Flexible Volume Metadata 16.23MB 0%
Delayed Frees 27.97MB 1%
Total Metadata Footprint 48.52MB 1%
Total Footprint 2.87GB 90%
Vserver : vs1
Volume : svm_root
Feature Used Used%
--- --- ---
Volume Data Footprint 2.93MB 0%
Footprint in Performance Tier 2.99MB 100%
Footprint in my-store 0B 0%
Volume Guarantee 17.01MB 0%
Flexible Volume Metadata 208KB 0%
Delayed Frees 68KB 0%
Total Metadata Footprint 320KB 0%
Total Footprint 20.20MB 0%
Vserver : vs1
Volume : vol1
Feature Used Used%
--- --- ---
Volume Data Footprint 1.61GB 17%
Footprint in Performance Tier 1.23GB 72%
Footprint in my-store 479.0MB 28%
Volume Guarantee 0B 0%
Flexible Volume Metadata 16.06MB 0%
Delayed Frees 82.98MB 1%
Total Metadata Footprint 103.52MB 1%
Total Footprint 1.71GB 18%
Vserver : vs1
Volume : vol2
Feature Used Used%

---------- ---------- -----
Volume Data Footprint 1.22GB 13%
Footprint in Performance Tier 823.3MB 65%
Footprint in ms1 440MB 35%
Volume Guarantee 0B 0%
Flexible Volume Metadata 16.06MB 0%
Delayed Frees 12MB 0%
Total Metadata Footprint 33.52MB 0%
Total Footprint 1.25GB 13%

4 entries were displayed.

```
== Related Links

* xref:{relative_path}volume-show-space.html[volume show-space]




[[IDe7b22d9121579e316d55067b8b6a30e3]]
= volume show-space
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display space usage for volume(s)

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume show-space` command displays information about space usage
within the volume. The command output depends on the parameter or
parameters specified with the command. If no parameters are specified, the
command displays the following information about all volumes.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter and the `-volume` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter and the `-vserver` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
all volumes matching the specified name.

``[-volume-msid <integer>]`` - Volume MSID::
If this parameter is specified, the command displays information only
about the volume that has the specified MSID.

``[-volume-dsid <integer>]`` - Volume DSID::
If this parameter is specified, the command displays information only
about the volume that has the specified DSID.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
If this parameter is specified, the command displays information only
about the volume on the vserver which has the specified UUID.

``[-aggregate <aggregate name>]`` - Aggregate Name::
If this parameter is specified, the command displays information only
about the volumes that are associated with the specified aggregate.

``[-aggregate-uuid <UUID>]`` - Aggregate UUID::
If this parameter is specified, the command displays information only
about the volumes on the aggregate which have the specified UUID.

``[-hostname <text>]`` - Hostname::
If this parameter is specified, the command displays information only
about the volumes that belong to the specified host.

``[-user-data {<integer>[KB|MB|GB|TB|PB]}]`` - User Data::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by user data in the volume.
+
This is the amount of data written to the volume via CIFS, NFS or SAN
protocols plus the space reserved in the volume for these files (hole and
overwrite reserves). This is the same information displayed by running the
```

```
Unix `du` command on the mount point.

``[-user-data-percent <percent_no_limit>]`` - User Data Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by user data in the volume.

``[-dedupe-metafiles {<integer>[KB|MB|GB|TB|PB]}]`` - Deduplication::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by deduplication metafiles in the volume.

``[-dedupe-metafiles-percent <percent>]`` - Deduplication Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by deduplication metafiles in the volume.

``[-dedupe-metafiles-temporary {<integer>[KB|MB|GB|TB|PB]}]`` - Temporary
Deduplication::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by temporary deduplication metafiles in the volume.

``[-dedupe-metafiles-temporary-percent <percent>]`` - Temporary
Deduplication Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by temporary deduplication metafiles in the volume.

``[-filesystem-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - Filesystem
Metadata::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by file system metadata in the volume.

``[-filesystem-metadata-percent <percent>]`` - Filesystem Metadata
Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by file system metadata in the volume.

``[-snapmirror-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - SnapMirror
Metadata::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by SnapMirror metafiles in the volume.
```

```
+
Between SnapMirror transfers, some metadata is maintained to support
storage-efficient transfers. During transfers, some additional space is
used temporarily. This space is used in all SnapMirror destination
volumes.

``[-snapmirror-metadata-percent <percent>]`` - SnapMirror Metadata
Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by SnapMirror metafiles inside the volume.

``[-tape-backup-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - Tape Backup
Metadata::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by tape backup metafiles in the volume.

``[-tape-backup-metadata-percent <percent>]`` - Tape Backup Metadata
Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by tape backup metafiles in the volume.

``[-quota-metafiles {<integer>[KB|MB|GB|TB|PB]}]`` - Quota Metadata::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by quota metafiles.

``[-quota-metafiles-percent <percent>]`` - Quota Metadata Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by quota metafiles.

``[-inodes {<integer>[KB|MB|GB|TB|PB]}]`` - Inodes::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by the inode metafile in the volume.
+
This is the amount of space required to store inodes in the file system
and is proportional to the maximum number of files ever created in the
volume. The inode file is not compacted or truncated, so if a large number
of files are created and then deleted, the inode file does not shrink.

``[-inodes-percent <percent>]`` - Inodes Percent::
If this parameter is specified, the command displays information only
```

```
about the volume or volumes that have the specified percentage of space in
use by the inode metafile in the volume.

``[-inodes-upgrade {<integer>[KB|MB|GB|TB|PB]}]`` - Inodes Upgrade::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by the inode subsystem for the purpose of upgrading.
+
This is the amount of space required to store upgrading inodes in the file
system and is proportional to the size of the inode metafile. Once the
upgrade is complete, the space used by 'inodes' will be replaced with the
space used for upgrade.

``[-inodes-upgrade-percent <percent>]`` - Inodes Upgrade Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use for upgrading the inode metafile in the volume.

``[-snapshot-reserve {<integer>[KB|MB|GB|TB|PB]}]`` - Snapshot Reserve::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by the snapshot reserve in the volume.

``[-snapshot-reserve-percent <percent>]`` - Snapshot Reserve Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by the snapshot reserve in the volume.

``[-snapshot-reserve-unusable {<integer>[KB|MB|GB|TB|PB]}]`` - Snapshot
Reserve Unusable::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space
reserved but unusable in the volume.
+
Snapshot reserve can be diminished under certain conditions to accommodate
volume metadata. Creating space in the volume will make this space
available.

``[-snapshot-reserve-unusable-percent <integer>]`` - Snapshot Reserve
Unusable Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space
reserved but unusable.

``[-snapshot-spill {<integer>[KB|MB|GB|TB|PB]}]`` - Snapshot Spill::
If this parameter is specified, the command displays information only
```

```
about the volume or volumes that have the specified amount of space in use
by their snapshot spill.
+
If snapshot used space exceeds the snapshot reserve it is considered to
spill out of the reserve. This space cannot be used by the active file
system until snapshots are deleted.

``[-snapshot-spill-percent <percent>]`` - Snapshot Spill Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by the snapshot spill.

``[-performance-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - Performance
Metadata::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
for performance optimization in the volume.

``[-performance-metadata-percent <percent>]`` - Performance Metadata
Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use for performance optimization in the volume.

``[-total-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total Used::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space in use
by the volume, including the space used by the snapshot reserve. ``_total-
used_`` differs from ``_physical-used_`` space by the sum of of the space
that is reserved for future writes and the space that is saved by
aggregate storage efficiency savings being removed from ``_total-used_`` .
+
This is equivalent to the used field in the output of the
xref:{relative_path}volume-show.html[volume show] command.

``[-total-used-percent <percent_no_limit>]`` - Total Used Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space in
use by the volume, including the space used by the snapshot reserve.

``[-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total Physical Used
Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of physical
space in use by the volume.
+
```

```
This value represents the effective total footprint which is the total
footprint after deducting auto adaptive compression and compaction
savings. For volumes with a `space-guarantee` of ``_volume_`` the value
also includes provisioned space. The value also includes aggregate space
used by the volume metadata. Because of this the value reported can be
more than 100%.
+
``_physical-used_`` differs from ``_total-used_`` space by the sum of the
space that is reserved for future writes and the space that is saved saved
by aggregate storage efficiency savings being added to ``_physical-used_``
. The value includes blocks in use by snapshots.

``[-physical-used-percent <percent_no_limit>]`` - Physical Used
Percentage::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of physical
space in use in the volume based on volume size including the space
reserved for snapshots.

``[-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Used Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of logical
space in use by the volume. This includes space saved by all the storage
efficiency features along with physical used space. This does not include
snapshot reserve but does consider snapshot spill.

``[-logical-used-percent <percent_no_limit>]`` - Logical Used Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of logical
space used in the volume.

``[-logical-available {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Available::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of logical
available space in the volume.

``[-total-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - Total Metadata::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space used
by metadata in the volume. This field will not include the metadata of the
volume which is residing inside the aggregate.

``[-total-metadata-percent <percent>]`` - Total Metadata Percentage::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space
used by metadata in the volume.
```

```
== Examples

The following example shows how to display details for all volumes.
```

cluster1::> volume show-space
Vserver : nodevs
Volume : vol0
Feature Used Used%
--- --- ---
User Data 163.4MB 3%
Filesystem Metadata 172KB 0%
Inodes 2.93MB 0%
Snapshot Reserve 292.9MB 5%
Total Metadata 185KB 0%
Total Used 459.4MB 8%
Total Physical Used 166.4MB 3%
Vserver : thevs
Volume : rootvol
Feature Used Used%
--- --- ---
User Data 100KB 0%
Filesystem Metadata 76KB 0%
Inodes 24KB 0%
Snapshot Reserve 1MB 5%
Total Metadata 89KB 0%
Total Used 1.20MB 6%
Total Physical Used 200KB 1%
Vserver : vs1
Volume : vol1
Feature Used Used%
--- --- ---
User Data 180.8MB 74%
Filesystem Metadata 280KB 0%
Inodes 12KB 0%
Snapshot Reserve 12.20MB 5%
Total Metadata 295KB 0%
Total Used 193.3MB 79%
Total Physical Used 192.9MB 79%
3 entries were displayed.

```
The following example shows all Volumes that have a snap reserve greater than 2 MB:
```

cluster1::> volume show-space -snapshot-reserve >2m
Vserver : nodevs
Volume : vol0
Feature Used Used%
--- --- ---
User Data 163.4MB 3%
Filesystem Metadata 172KB 0%
Inodes 2.93MB 0%

Snapshot Reserve 292.9MB 5%
Total Metadata 3.21MB 0%
Total Used 459.4MB 8%
Total Physical Used 166.4MB 3%
Vserver : vs1
Volume : vol1
Feature Used Used%
--- --- ---
User Data 180.8MB 74%
Filesystem Metadata 280KB 0%
Inodes 12KB 0%
Snapshot Reserve 12.20MB 5%
Total Metadata 312KB 0%
Total Used 193.3MB 79%
Total Physical Used 192.9MB 79%
2 entries were displayed.

```
== Related Links

* xref:{relative_path}volume-show.html[volume show]




[[ID965bfaf2fa8a2c5febabc5531a8d8032]]
= volume show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume show` command displays information about volumes. The command
output depends on the parameter or parameters specified with the command.
If no parameters are specified, the command displays the following
information about all volumes:

* Vserver name
* Volume name
* Aggregate name
* State (online, offline, restricted, or mixed)
```

```
* Type (RW for read-write or DP for data-protection)
* Size
* Available size
* Percentage of space used



To display detailed information about a single volume, run the command
with the `-vserver` and `-volume` parameters. The detailed view provides
all of the information in the previous list and the following additional
information:

* Name ordinal
* Volume data set ID
* Volume master data set ID
* Volume style (trad or flex)
* FlexCache Endpoint Type (none or cache or origin)
* Whether the volume is a cluster volume or node volume
* Export policy name
* User ID
* Group ID
* Security style (unix, ntfs, mixed or unified)
* UNIX permissions
* Junction path
* Junction path source
* Whether the junction path is active
* Parent volume name
* Vserver root volume
* Comment
* Filesystem size
* Total user-visible size
* Used size
* Used percentage
* Volume nearly full threshold percent
* Volume full threshold percent
* Autosize enabled
* Maximum autosize
* Minimum autosize
* Autosize grow threshold percent
* Autosize shrink threshold percent
* Autosize mode
* Total files
* Files used
* Expected available size
* Over provisioned size
* Snapshot reserve available size
```

```
* Logical used size
* Logical used percent
* Logical available size
* Active filesystem logical used size
* Snapshot logical used size
* Use logical space reporting
* Use logical space enforcement
* Maximum directory size
* Space guarantee style
* Whether a space guarantee is in effect
* Space SLO type (none, thick or semi-thick)
* Whether space SLO is in effect
* Whether minimum readahead is enabled
* Whether access time update is enabled
* Whether snapshot directory access is enabled
* Percentage of space reserved for snapshots
* Percentage of snapshot space used
* Snapshot policy name
* Creation time
* If the filesystem size is fixed
* Overwrite reserve
* Fractional reserve
* Which space management strategy to try first
* Language
* Whether there's one data volume per member aggregate
* Concurrency level
* Optimization policy
* Whether the volume is a clone
* Volume UUID
* Whether failover is enabled
* Failover state
* (DEPRECATED)-Extent option
* Read reallocation option
* Consistency state
* Whether volume is quiesced on disk
* Whether volume is quiesced in memory
* Whether volume contains shared or compressed data
* Space saved by storage efficiency
* Percentage of space saved by storage efficiency
* Space saved by deduplication
* Percentage of space saved by deduplication
* Space shared by deduplication
* Space saved by compression
* Percentage of space saved by compression
* Volume size used by snapshots
* Caching policy
```

```
* FlexGroup volume master data set ID
* FlexGroup volume index
* FlexGroup volume UUID
* Maximum size of the FlexGroup volume constituent
* Whether the volume has FlexGroup volume enabled
* Whether a FlexGroup volume is Qtree enabled
* Whether the volume is the destination of a move that is currently in
cutover
* List of the aggregates used by the FlexGroup volume
* List of the nodes used by the FlexGroup volume
* SnapLock Type
* Is in pre-commit phase of copy-free Transition
* Application that the volume belongs to
* File system analytics state of the volume
* File system analytics scan progress percentage
* File system analytics number of files scanned
* File system analytics total number of files
* Activity tracking state of the volume
* Application UUID
* Whether the volume's total number of files is set to the highest
possible value

To display detailed information about all volumes, run the command with
the `-instance` parameter.

You can specify additional parameters to display information that matches
only those parameters. For example, to display information only about
data-protection volumes, run the command with the `-type DP` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed. The fields Vserver
and policy are the default fields (see example).

| ``[-encryption  ]``::
If this parameter is specified, the command displays the following
information:
+
  * Vserver name
  * Volume name
  * Aggregate name
  * Volume state
  * Encryption state
```

```
| ``[-junction  ]``::
If this parameter is specified, the command displays the following
information:
+
  * Vserver name
  * Volume name
  * Whether the volume's junction is active
  * Junction path
  * Junction path source (if the volume is a mirror)



| ``[-settings  ]`` [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the following
information:
+
  * Vserver name
  * Volume name
  * Whether minimum readahead is enabled on the volume
  * Whether the access time is updated on inodes when a file on the volume
is read
  * Whether clients have access to .snapshot directories
  * Whether automatic snapshots are enabled on the volume



| ``[-instance  ]`` }::
If this parameter is specified, the command displays information about all
entries.

``[-vserver <vserver name>]`` - Vserver Name::
If this parameter and the `-volume` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter and the `-vserver` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
all volumes matching the specified name.

``[-aggregate <aggregate name>]`` - Aggregate Name::
If this parameter is specified, the command displays information only
about the volume or volumes that are located on the specified storage
aggregate. This field is displayed as "-" for FlexGroup volumes.
```

```
``[-aggr-list <aggregate name>,...]`` - List of Aggregates for FlexGroup
Constituents::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volumes that are located on the
specified list of storage aggregates. This parameter applies to FlexGroup
volumes only.

``[-encryption-type {none|volume|aggregate|mixed}]`` - Encryption Type::
If this parameter is specified, the command displays information about the
type of encryption key used for encrypting the volume. The possible values
are ``_none_`` , ``_volume_`` , and ``_aggregate_`` . The value ``_none_``
is used for non-encrypted volumes, the value ``_volume_`` is used for
volumes encrypted with NVE (NetApp Volume Encryption) and ``_aggregate_``
is used for volumes encrypted with NAE (NetApp Aggregate Encryption).

``[-nodes {<nodename>|local}]`` - List of Nodes Hosting the Volume::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volumes that are located on the
specified list of storage systems. This parameter applies to FlexGroup
volumes only.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified size. Size is the
maximum amount of space a volume can consume from its associated
aggregate(s), including user data, metadata, snapshots, and snapshot
reserve. Note that for volumes without a `-space-guarantee` of
``_volume_`` , the ability to fill the volume to this maximum size depends
on the space available in the associated aggregate or aggregates.

``[-name-ordinal <text>]`` - Name Ordinal [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, it denotes the ordinal assignment used in
relation to this volume's name. Ordinals are used to disambiguate volumes
that have the same base name on the same controller. A value of "0"
indicates that the base volume name is unique on the controller. A value
greater than zero indicates that the volume's base name is used by two or
more volumes on the same controller, and that appending "(n)" to this
volume's name uniquely identifies it on this controller.

``[-dsid <integer>]`` - Volume Data Set ID::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified data set ID. This
field is displayed as "-" for FlexGroup volumes.

``[-msid <integer>]`` - Volume Master Data Set ID::
```

```
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified master data set ID.

``[-state {online|restricted|offline|force-online|force-offline|mixed}]``
- Volume State::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified state. The ``_mixed_``
state only applies to FlexGroup volumes. If the state of a FlexGroup
volume is ``_mixed_`` , that indicates that not all of the constituents
are in the same state. If this is the case use the "volume show -is
-constituent true" command to find out which constituents are not in the
proper state.

``[-volume-style <flex>]`` - Volume Style::
If this parameter is specified, the command displays information only
about the volumes that have the specified style. Possible values are
``_flex_`` for FlexVol volumes.

``[-volume-style-extended {flexvol|flexgroup|flexgroup-constituent}]`` -
Extended Volume Style::
If this parameter is specified, the command displays information only
about the volumes that are configured with the specified extended style.
Possible values are ``_flexvol_`` for FlexVol volumes, ``_flexgroup_`` for
FlexGroup volumes and ``_flexgroup-constituent_`` for FlexGroup volume
constituents.

``[-flexcache-endpoint-type {none|cache|origin}]`` - FlexCache Endpoint
Type::
If this parameter is specified, the command displays information only
about the volumes that are of the specified flexcache-endpoint-type.
Possible values are ``_none_`` for volumes that are not part of a
FlexCache relationship, ``_cache_`` for FlexCache volumes and ``_origin_``
for origin of FlexCache volumes.

``[-is-cluster-volume {true|false}]`` - Is Cluster-Mode Volume::
If this parameter is specified, the command displays information only
about cluster volumes (true) or node root volumes and other node scoped
volumes (false).

``[-is-constituent {true|false}]`` - Is Constituent Volume::
If this parameter is specified, the command displays information only
about volumes that either are or are not constituents of a FlexGroup
volume, depending on the value provided.

``[-constituent-count <integer>]`` - Number of Constituent Volumes::
The number of constituents in the FlexGroup volume. This parameter applies
```

```
to FlexGroup volumes only.

``[-policy <text>]`` - Export Policy::
If this parameter is specified, the command displays information only
about the volume or volumes that use the specified export policy.

``[-user <user name>]`` - User ID::
If this parameter is specified, the command displays information only
about the volume or volumes whose root is owned by the specified user.

``[-group <group name>]`` - Group ID::
If this parameter is specified, the command displays information only
about the volume or volumes whose root is owned by the specified group.

``[-security-style <security style>]`` - Security Style::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified security style
(``_unix_`` for UNIX mode bits, ``_ntfs_`` for CIFS ACLs, ``_mixed_`` for
both styles or ``_unified_`` for Unified UNIX, NFS and CIFS permissions).

``[-unix-permissions <unix perm>]`` - UNIX Permissions::
If this parameter is specified, the command displays information only
about the volume or volumes whose default UNIX permissions match the
specified permissions. Specify UNIX permissions either as a four-digit
octal value (for example, 0700) or in the style of the UNIX ls command
(for example, -rwxr-x---). For information on UNIX permissions, see the
UNIX or Linux documentation.

``[-junction-path <junction path>]`` - Junction Path::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified junction path.

``[-junction-path-source {RW_volume|LS_mirror}]`` - Junction Path Source::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified junction path source.

``[-junction-active {true|false}]`` - Junction Active::
If this parameter is specified, the command displays information only
about the volume or volumes whose junction paths have the specified
status.

``[-junction-parent <volume name>]`` - Junction Parent Volume::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified parent volume.

``[-vsroot {true|false}]`` - Vserver Root Volume [.advaced-
priv]#(privilege: advanced)#::
```

```
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified setting; that is,
whether they are the root volumes for their Vservers.

``[-comment <text>]`` - Comment::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified comment text.

``[-available {<integer>[KB|MB|GB|TB|PB]}]`` - Available Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified available size.
Available is the amount of free space currently available to be used by
this volume. For a volume with a `-space-guarantee` of type ``_volume_`` ,
available is always `-total` minus `-used` . For volumes that do not have
a `-space-guarantee` of type ``_volume_`` , available could be reduced if
the volume's associated aggregate or aggregates are space constrained.

``[-filesystem-size {<integer>[KB|MB|GB|TB|PB]}]`` - Filesystem Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified filesystem size.
Filesystem size is the same as the volume's `-size` unless the volume is
or was a physical replica destination. In this case, the file system size
corresponds to the `-size` of the source volume, until `-filesys-size-
fixed` is set to ``_false_`` .

``[-total {<integer>[KB|MB|GB|TB|PB]}]`` - Total User-Visible Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified total size. Total is
the total space available for user data and file system metadata. It does
not include the snapshot reserve.

``[-used {<integer>[KB|MB|GB|TB|PB]}]`` - Used Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified used size. Used is the
amount of space occupied by user data and file system metadata. ``_used_``
differs from ``_physical-used_`` space by the sum of the space that is
reserved forfeiture writes and the space that is saved by aggregate
storage efficiency savings being removed from ``_used_`` . It includes
snapshot spill (the amount of space by which snapshots exceed snapshot
reserve). It does not include the snapshot reserve.

``[-percent-used <percent>]`` - Used Percentage::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of used
space. This row is based on a value of used space that includes the space
used by snapshots or the snapshot reserve (whichever is greater) in
```

```
relation to the current volume size.

``[-space-nearly-full-threshold-percent <percent>]`` - Volume Nearly Full
Threshold Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified nearly full threshold
percent.

``[-space-full-threshold-percent <percent>]`` - Volume Full Threshold
Percent::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified full threshold
percent.

``[-max-autosize {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Autosize::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified maximum automatic
size.

``[-min-autosize {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Autosize::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified minimum automatic
size. This field is displayed as "-" for FlexGroup volumes.

``[-autosize-grow-threshold-percent <percent>]`` - Autosize Grow Threshold
Percentage::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified automatic grow used
space threshold percentage. This field is displayed as "-" for FlexGroup
volumes.

``[-autosize-shrink-threshold-percent <percent>]`` - Autosize Shrink
Threshold Percentage::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified automatic shrink used
space threshold percentage. This field is displayed as "-" for FlexGroup
volumes.

``[-autosize-mode {off|grow|grow_shrink}]`` - Autosize Mode::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified automatic sizing mode
setting. This field is displayed as "-" for FlexGroup volumes.

``[-files <integer>]`` - Total User Visible Files::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of files.
```

```
``[-files-used <integer>]`` - User Visible Files Used::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of files used.
This includes all user-visible file system entities such as directories,
regular files, streams, ACLS, etc.

``[-inodefile-public-capacity <integer>]`` - User Visible Inode File
Capacity [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified public inode file
capacity. This includes all user-invisible or metadata files such as
private metafiles and zombies.

``[-files-maximum-possible <integer>]`` - Maximum Possible Total User
Visible Files [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of maximum
files possible. This includes all user-visible file system entities such
as directories, regular files, streams, ACLS, etc.

``[-files-set-maximum {true|false}]`` - Set Total Files (for user-visible
data) to the Highest Value that the Volume can Hold [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information about
whether the volume's total number of files is set to the highest value
that the volume can hold. This includes all user-visible file system
entities such as directories, regular files, streams, ACLS, etc.

``[-files-private-used <integer>]`` - Private System Metadata Files Used
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of private
system metadata files used. This includes all user-invisible or metadata
files such as private metafiles and zombies.

``[-inodefile-private-capacity <integer>]`` - Private System Metadata
Inode File Capacity [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified private inode file
capacity. This includes all user-visible file system entities such as
directories, regular files, streams, ACLS, etc.

``[-maxdir-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Directory Size
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified maximum directory
```

```
size.

``[-space-guarantee-enabled {true|false}]`` - Space Guarantee in Effect::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified space-guarantee
setting. If the value of `-space-guarantee` is ``_none_`` , the value of
`-space-guarantee-enabled` is always ``_true_`` . In other words, because
there is no guarantee, the guarantee is always in effect. If the value of
`-space-guarantee` is ``_volume_`` , the value of `-space-guarantee-
enabled` can be ``_true_`` or ``_false_`` , depending on whether the
guaranteed amount of space was available when the volume was mounted.

``[-is-space-slo-enabled {true|false}]`` - Space SLO in Effect::
If this parameter is specified, the command displays information only
about the volume or volumes that have their `space-slo` setting in effect
or not, depending on the value specified for this parameter. If the value
of `space-slo` is ``_none_`` , the space SLO is always considered to be in
effect. If the value of `space-slo` is ``_semi-thick_`` or ``_thick_`` ,
the space SLO may be in effect depending on whether the required amount of
space was available when the volume was mounted.

``[-space-slo {none|thick|semi-thick}]`` - Space SLO::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified `space-slo` setting.
The space SLO setting is the Service Level Objective for space management
for the volume.

``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified space guarantee style.
If the value of `-space-guarantee` is ``_none_`` , the value of `-space-
guarantee-enabled` is always ``_true_`` . In other words, because there is
no guarantee, the guarantee is always in effect. If the value of `-space-
guarantee` is ``_volume_`` , the value of `-space-guarantee-enabled` can
be ``_true_`` or ``_false_`` , depending on whether the guaranteed amount
of space was available when the volume was mounted.

``[-fractional-reserve <percent>]`` - Fractional Reserve::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified `fractional-reserve`
setting.

``[-type {RW|DP}]`` - Volume Type::
If this parameter is specified, the command displays information only
about the volume or volumes of the specified volume type (RW for read-
write or DP for data-protection).
```

```
``[-min-readahead {true|false}]`` - Minimum Read Ahead [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified minimum-readahead
setting.

``[-atime-update {true|false}]`` - Access Time Update Enabled [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified access-time update
setting.

``[-snapdir-access {true|false}]`` - Snapshot Directory Access Enabled::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified snapshot directory
access setting.

``[-percent-snapshot-space <percent>]`` - Space Reserved for Snapshot
Copies::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified percentage of space
reserved for snapshots.

``[-snapshot-space-used <percent_no_limit>]`` - Snapshot Reserve Used::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified used percentage of the
reserve for snapshots.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
If this parameter is specified, the command displays information only
about the volume or volumes that use the specified snapshot policy.

``[-create-time <Date>]`` - Creation Time::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified creation time.

``[-language <Language code>]`` - Language::
If this parameter is specified, the command displays information only
about the volume or volumes that store data in the specified language. To
determine the available languages, enter `volume show``-language```_?_``
at the clustershell command prompt.

``[-clone-volume {true|false}]`` - Clone Volume::
If this parameter is specified, the command displays information only
about volumes that are clones (true) or not clones (false).

``[-node {<nodename>|local}]`` - Node name::
```

```
If this parameter is specified, the command displays information only the
volume or volumes that are located on the specified storage system. This
field is displayed as "-" for FlexGroup volumes.

``[-clone-parent-vserver <vserver name>]`` - Clone Parent Vserver Name::
If this parameter is specified, the command displays information only
about the volumes with a matching FlexClone parent Vserver name.

``[-clone-parent-name <volume name>]`` - FlexClone Parent Volume::
If this parameter is specified, the command displays information only
about the volumes with a matching FlexClone parent volume name.

``[-uuid <UUID>]`` - UUID of the Volume [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified UUID.

``[-nvfail {on|off}]`` - NVFAIL Option::
If this parameter is specified, the command displays information only
about volumes for which failover is enabled (on) or disabled (off).

``[-in-nvfailed-state {true|false}]`` - Volume's NVFAIL State::
If this parameter is specified, the command displays information only
about volumes which are in the failed over state (true) or not (false).
This field is only available when the volume is online.

``[-dr-force-nvfail {on|off}]`` - Force NVFAIL on MetroCluster
Switchover::
If this parameter is specified, the command displays information only
about volumes for which dr-force-nvfail is enabled (on) or disabled (off).

``[-filesys-size-fixed {true|false}]`` - Is File System Size Fixed::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified `filesys-size-fixed`
setting.

``[-extent-enabled {off|on|space-optimized}]`` - (DEPRECATED)-Extent
Option::
+
NOTE: This parameter has been deprecated and might be removed in a future
release of ONTAP.
+
If this parameter is specified, the command displays information only
about volumes that have extents enabled (on), not enabled (off) or space
optimized (space-optimized).

``[-overwrite-reserve {<integer>[KB|MB|GB|TB|PB]}]`` - Reserved Space for
```

```
Overwrites::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified `overwrite-reserve`
setting.

``[-space-mgmt-try-first {volume_grow|snap_delete}]`` - Primary Space
Management Strategy::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified `space-mgmt-try-first`
setting. Possible values are `volume_grow` and `snap_delete` . This field
is displayed as "-" for FlexGroup volumes.

``[-read-realloc {off|on|space-optimized}]`` - Read Reallocation Option::
If this parameter is specified, the command displays information only
about volumes that have read reallocation enabled (on), not enabled (off)
or space optimized (space-optimized).

``[-sched-snap-name {create-time|ordinal}]`` - Naming Scheme for Automatic
Snapshot Copies::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified automatic snapshot
naming convention.

``[-is-inconsistent {true|false}]`` - Inconsistency in the File System::
If this parameter is specified, the command displays information only
about volumes that are inconsistent (true) or consistent (false) in the
file system.

``[-is-quiesced-on-disk {true|false}]`` - Is Volume Quiesced (On-Disk)::
If this parameter is specified, the command displays information only
about volumes that are quiesced (true) or not quiesced (false) on disk.

``[-is-quiesced-in-memory {true|false}]`` - Is Volume Quiesced (In-
Memory)::
If this parameter is specified, the command displays information only
about volumes that are quiesced (true) or not quiesced (false) in memory.

``[-transition-state <state>]`` - Transition Operation State [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified transition state.

``[-transition-behavior {data-move|data-protection|none}]`` - Transition
Behavior [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified transition behavior.
```

```
Possible values are:
+
  *  data-move: Volumes that are being moved from a system operating in 7-
Mode.
  *  data-protection: Volumes that are being replicated from a system
operating in 7-Mode for disaster recovery.
  *  none: Volumes that are not part of transition.



``[-is-copied-for-transition {true|false}]`` - Copied for Transition
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified value based on
whether the volume is copied for transition or not.

``[-is-transitioned {true|false}]`` - Transitioned [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified value based on
whether the volume is transitioned or not.

``[-is-sis-volume {true|false}]`` - Volume Contains Shared or Compressed
Data::
If this parameter is specified, the command displays information only
about those volumes that match the specified storage efficiency setting.

``[-sis-space-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved by
Storage Efficiency::
If this parameter is specified, the command displays information only
about those volumes that have the specified amount of space saved by the
storage efficiency technology.

``[-sis-space-saved-percent <percent>]`` - Percentage Saved by Storage
Efficiency::
If this parameter is specified, the command displays information only
about those volumes that have the specified percentage of space saved by
the storage efficiency technology.

``[-dedupe-space-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved by
Deduplication Along With VBN ZERO Savings::
If this parameter is specified, the command displays information only
about those volumes that have the specified amount of space saved due to
deduplication.

``[-dedupe-space-saved-percent <percent>]`` - Percentage Saved by
```

```
Deduplication::
If this parameter is specified, the command displays information only
about those volumes that have the specified percentage of space saved due
to deduplication.

``[-dedupe-space-shared {<integer>[KB|MB|GB|TB|PB]}]`` - Unique Data Which
Got Shared by Deduplication::
If this parameter is specified, the command displays information only
about those volumes that have the specified amount of shared space due to
deduplication.

``[-compression-space-saved {<integer>[KB|MB|GB|TB|PB]}]`` - Space Saved
by Compression::
If this parameter is specified, the command displays information only
about those volumes that have the specified amount of space saved due to
compression.

``[-compression-space-saved-percent <percent>]`` - Percentage Space Saved
by Compression::
If this parameter is specified, the command displays information only
about those volumes that have the specified percentage of space saved due
to compression.

``[-size-used-by-snapshots {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Size
Used by Snapshot Copies::
If this parameter is specified, the command displays information about
those volumes that have the specified volume size used by snapshots.

``[-block-type {64-bit|extent|32-bit}]`` - Block Type::
If this parameter is specified, the command displays information about
only the volumes with the specified indirect block format. Possible values
are ``_32-bit_`` to display 32-bit volumes and ``_64-bit_`` to display 64-
bit volumes.

``[-is-moving {true|false}]`` - Is Volume Moving::
If this parameter is specified, the command displays information only
about volumes that are moving (true) or not moving (false).

``[-hybrid-cache-eligibility {read|read-write|none}]`` - Flash Pool
Caching Eligibility::
If this parameter is specified, the command displays information only
about the volume or volumes with the specified Flash Pool caching
attributes. Possible caching attributes are:
+
  *  'read' ... Indicates that the volume cannot participate in write
caching.
```

```
  *  'read-write' ... Indicates that the volume can participate in read
and write caching.



``[-hybrid-cache-write-caching-ineligibility-reason <text>]`` - Flash Pool
Write Caching Ineligibility Reason::
If this parameter is specified, the command displays information only
about the volume or volumes which are ineligible to participate in write
caching due to the specified reason.

``[-constituent-role <Constituent Roles>]`` - Constituent Volume Role::
If this parameter is specified, the command displays information only
about the constituent volume or volumes that are of the specified role.
This parameter applies to FlexGroup volumes only.

``[-is-cft-precommit {true|false}]`` - Is in the precommit phase of Copy-
Free Transition [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified with the true value, it displays
information only about the volumes that are in the precommit phase of a
Copy-Free Transition workflow.

``[-qos-policy-group <text>]`` - QoS Policy Group Name::
If this parameter is specified, the command displays information only
about volumes that match the specified QoS policy group.

``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name::
If this parameter is specified, the command displays information only
about volumes that match the specified QoS adaptive policy group.

``[-caching-policy <text>]`` - Caching Policy Name::
 If this parameter is specified, the command displays the volumes that
match the specified caching policy.
+
A caching policy defines how the system caches a volume's data in a Flash
Pool aggregate. Both metada and user data are eligible for caching. The
available caching policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
```

```
  * all_read - Read caches all metadata, randomly read and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read, and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.
  * noread-random_write - Write caches all randomly overwritten user data
blocks. It does not do any read caching.
  * meta-random_write - Read caches all metadata and write caches randomly
overwritten user data blocks.
  * random_read_write-random_write - Read caches all metadata, randomly
read and randomly written user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read-random_write - Read caches all metadata, randomly read and
sequentially read user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read_random_write-random_write - Read caches all metadata,
randomly read, sequentially read, and randomly written user data. It also
write caches randomly overwritten user data blocks.
  * all-random_write - Read caches all data blocks read and written. It
also write caches randomly overwritten user data blocks.

+
Note that in a caching-policy name, a hyphen (-) separates read and write
policies. Default caching-policy is auto.

``[-cache-retention-priority {normal|low|high}]`` - Cache Retention
Priority [.advaced-priv]#(privilege: advanced)#::
 If this parameter is specified, the command displays the volumes that
match the specified cache retention priority policy.
+
A cache retention priority defines how long the blocks of a volume will be
cached in Flash Pool once they become cold.The available cache retention
priority are:
+
  * low - Cache the cold blocks for the lowest time.
  * normal - Cache the cold blocks for the default time.
  * high - Cache the cold blocks for the highest time.



``[-is-volume-in-cutover {true|false}]`` - Is Volume Move in Cutover
Phase::
If this parameter is specified, the command displays information only
about volumes that are in the cutover phase (true) or not in the cutover
phase (false) of a volume move. This field is displayed as "-" for
```

```
FlexGroup volumes.

``[-snapshot-count <integer>]`` - Number of Snapshot Copies in the
Volume::
If this parameter is specified, the command displays information only
about the volumes that have the specified number of snapshots.

``[-vbn-bad-present {true|false}]`` - VBN_BAD may be present in the active
filesystem::
If this parameter is specified, the command displays information only
about volumes that may have VBN_BAD present in its active filesystem
(true) or do not have VBN_BAD present in its active filesystem (false).

``[-is-autobalance-eligible {true|false}]`` - Is Eligible for Auto Balance
Aggregate [.advaced-priv]#(privilege: advanced)#::
 If this parameter is specified, the command displays information only
about the volume or volumes that are eligible for consideration by the
Auto Balance Aggregate feature.

``[-is-vol-on-hybrid-aggr {true|false}]`` - Is Volume on a hybrid
aggregate::
If this parameter is specified, the command displays information only
about volumes associated with a Flash Pool aggregate (true) or not
(false). This field is displayed as "-" for FlexGroup volumes.

``[-physical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Total Physical Used
Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified physical used size.
Physical-used represents effective total footprint which is total
footprint after deducting auto adaptive compression and compaction
savings. For volumes with a `space-guarantee` of ``_volume_`` the value
includes provisioned space. The value also includes aggregate space used
by the volume metadata. Because of this the value reported can be more
than 100%. ``_physical-used_`` differs from ``_total-used_`` space by the
sum of the space that is reserved for future writes and the space that is
saved by aggregate storage efficiency savings being added to ``_physical-
used_`` . The value includes blocks in use by snapshots.

``[-physical-used-percent <percent_no_limit>]`` - Physical Used
Percentage::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified physical used percent
based on volume size including the space reserved for snapshots.

``[-flexgroup-msid <integer>]`` - FlexGroup Master Data Set ID [.advaced-
```

```
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volume constituents that have the
specified FlexGroup volume master data-set ID. This parameter applies to
FlexGroup volumes and FlexGroup volume constituents only.

``[-flexgroup-index <integer>]`` - FlexGroup Index [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexGroup volume constituents that have the specified FlexGroup
volume index. This parameter applies to FlexGroup volume constituents
only.

``[-flexgroup-uuid <UUID>]`` - UUID of the FlexGroup [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volume constituents that have the
specified FlexGroup volume UUID. This parameter applies to FlexGroup
volumes and FlexGroup volume constituents only.

``[-flexgroup-name <volume name>]`` - FlexGroup Name::
If this parameter is specified, the command displays information only
about the volumes that have the specified FlexGroup volume name. This
parameter only applies to FlexGroup volumes and FlexGroup volume
constituents.

``[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum size of
a FlexGroup Constituent [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volumes that have the specified
maximum constituent size. This parameter applies to FlexGroup volumes
only.

``[-inofile-version <integer>]`` - Inofile Version [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volumes whose inode files are at the specified version.

``[-is-flexgroup {true|false}]`` - Is Volume a FlexGroup::
If this parameter is specified, the command displays information only
about the volume or volumes that are either FlexGroup volumes or not,
depending on the value provided.

``[-is-qtree-caching-enabled {true|false}]`` - Is Qtree Caching Support
Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
```

```
about the FlexGroup volumes and origins of FlexCache volumes with Qtree
caching enabled or disabled, depending on the value provided.

``[-is-move-destination-in-cutover {true|false}]`` - Is the Volume a
Target of a Move Which Is Currently in Cutover [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays whether or not the
volume is a move destination that is currently in cutover.

``[-snaplock-type {non-snaplock|compliance|enterprise}]`` - SnapLock
Type::
If this parameter is specified, the command displays information only
about volumes that match the specified snaplock-type.

``[-vserver-dr-protection {protected|unprotected}]`` - Vserver DR
Protection::
 If this parameter is specified, the command displays information only
about the volumes having the specified Vserver SnapMirror protection.

``[-has-optimized-sparse-directories {true|false}]`` - Volume Has Sparse
Directories in an Optimized Format [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about volumes that have sparse directories in an optimized format.

``[-encrypt {true|false}]`` - Enable or Disable Encryption::
If this parameter is specified, the command displays information only
about the volumes that are encrypted (true) or not encrypted (false).

``[-is-encrypted {true|false}]`` - Is Volume Encrypted::
If this parameter is specified, the command displays information only
about the volumes that are encrypted (true) or unencrypted (false).

``[-encryption-state
{none|full|partial|converting_to_encrypted|converting_to_plaintext|rekeyin
g|conversion_paused|rekey_paused}]`` - Encryption State::
If this parameter is specified, the command displays information only
about the volumes that have the specified encryption state. The possible
values are ``_none_`` , ``_full_`` , ``_partial_`` ,
``_converting_to_encrypted_`` , ``_converting_to_plaintext_`` , and
``_rekeying_`` . The value ``_partial_`` is used for FlexGroup volumes,
which indicates that some constituents are encrypted and some are not.

``[-key-id <text>]`` - Encryption Key ID::
If this parameter is specified, the command displays information only
about the volume whose encryption key-id matches the specified key-id.
```

```
``[-key-creation-time <MM/DD/YYYY HH:MM:SS>]`` - Encryption Key Creation
Time::
If this parameter is specified, the command displays information only
about volumes having an encryption key with a creation time that matches,
is greater than or less than the provided key-creation-time value. For
example, the `volume show -key-creation-time >1/1/2022 13:12:11` command
will show all volumes that have an encryption key created later than
1/1/2022 13:12:11

``[-application <text>]`` - Application::
Selects the volumes that are part of an application that matches the
parameter value.

``[-is-protocol-access-fenced {true|false}]`` - Is Fenced for Protocol
Access::
If this parameter is specified, the command displays information only
about the volumes that are fenced for protocol access. Only FlexGroup
volume constituents and volumes in SnapMirror Synchronous relationships
can be fenced for protocol access.

``[-protocol-access-fenced-by
{none|coordinated_snaprestore|coordinated_redirection|snapmirror_synchrono
us|vserver_migrate}]`` - Protocol Access Fence Owner::
This field indicates the owner of the protocol access fence when the
volume's protocol access is fenced. Only FlexGroup volume constituents and
volumes in SnapMirror Synchronous relationships can be fenced for protocol
access.

``[-single-instance-data-logging {off|on}]`` - Is SIDL enabled::
If this parameter is specified, the command displays whether Single
Instance Data Logging feature is enabled on the specified volume.

``[-over-provisioned {<integer>[KB|MB|GB|TB|PB]}]`` - Over Provisioned
Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified size unavailable in
the aggregate. This applies only for a "none" guaranteed volume when its
unused size exceeds available space in the aggregate. This value is always
zero for "volume" guaranteed volumes. For FlexGroup volumes, the value is
not the direct sum of the constituents’ overprovisioned size. Instead the
value represents the total overprovisioned space across all aggregates for
the FlexGroup volume. For each aggregate, the value is computed as the sum
of the unused size of the thin provisioned constituents of the FlexGroup
volumes that exceed available space in the aggregate.

``[-snapshot-reserve-available {<integer>[KB|MB|GB|TB|PB]}]`` - Available
```

```
Snapshot Reserve Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified size available for
snapshots within the snapshot reserve. This value is zero if snapshot
spill is present. For 'none' guaranteed volumes, this may get reduced due
to less available space in the aggregate.

``[-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Used Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified logical used size.
This value includes all the space saved by the storage efficiency features
along with the physically used space. This does not include snapshot
reserve but does consider snapshot spill.

``[-logical-used-percent <percent_no_limit>]`` - Logical Used Percentage::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified logical used
percentage.

``[-logical-available {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Available
Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified logical available
size. This value is the amount of free space currently available
considering space saved by the storage efficiency features as being used.
This does not include snapshot reserve.

``[-logical-used-by-afs {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Size Used
by Active Filesystem::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified logical size used by
the active file system. This value differs from ``_logical-used_`` by the
amount of snapshot spill that exceeds snapshot reserve. This parameter is
not supported on FlexGroup volumes.

``[-logical-used-by-snapshots {<integer>[KB|MB|GB|TB|PB]}]`` - Logical
Size Used by All Snapshots::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified logical size used
across all snapshots. This value differs from ``_size-used-by-snapshots_``
by the space saved by the storage efficiency features across the
snapshots. This parameter is not supported on FlexGroup volumes.

``[-is-space-reporting-logical {true|false}]`` - Logical Space Reporting::
If this parameter is specified, the command displays information only
about the volumes that have logical space reporting enabled or disabled as
```

```
specified. When space is reported logically, ONTAP reports the volume
space such that all the physical space saved by the storage efficiency
features are also as reported as used.

``[-is-space-enforcement-logical {true|false}]`` - Logical Space
Enforcement::
If this parameter is specified, the command displays information only
about the volumes that have logical space enforcement enabled or disabled
as specified. When space is enforced logically, ONTAP enforces volume
settings such that all the physical space saved by the storage efficiency
features will be calculated as used.

``[-tiering-policy <Tiering Policy>]`` - Volume Tiering Policy::
If this parameter is specified, the command displays information only
about the volumes whose tiering policy matches the specified value.
Tiering policies determine whether the user data blocks of a volume in a
FabricPool will be tiered to the cloud tier when they become cold.
FabricPool combines Flash (performance tier) with an object store (cloud
tier) into a single aggregate. The temperature of a volume block increases
if it is accessed frequently and decreases when it is not.
+
The available tiering policies are:
+
  * snapshot-only - Only the volume snapshots not associated with the
active file system are tiered to the cloud tier.
  * auto - Both snapshot data and active file system user data are tiered
to the cloud tier.
  * none - No volume blocks are tiered to the cloud tier.
  * all - Both snapshot data and active file system user data are tiered
to the cloud tier as soon as possible without waiting for a cooling
period. On DP volumes all transferred user data blocks start in the cloud
tier.



``[-cloud-retrieval-policy {default|on-read|never|promote}]`` - Volume
Cloud Retrieval Policy [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the cloud retrieval policy for the
volume. This policy determines which tiered out blocks to retrieve from
the capacity tier to the performance tier.
+
The available cloud retrieval policies are:
+
  * default - This policy retrieves tiered data based on the underlying
tiering policy. If the tiering policy is 'auto', tiered data is retrieved
only for random client driven data reads. If the tiering policy is 'none'
```

```
or 'snapshot-only', tiered data is retrieved for random and sequential
client driven data reads. If the tiering policy is 'all', tiered data is
not retrieved.
  * on-read - This policy retrieves tiered data for all client driven data
reads.
  * never - This policy never retrieves tiered data.
  * promote - This policy retrieves all eligible tiered data automatically
during the next scheduled scan. It is only supported when the tiering
policy is 'none' or 'snapshot-only'. If the tiering policy is 'snapshot-
only', the only data brought back is the data in the AFS. Data that is
only in a snapshot stays in the cloud.



``[-tiering-minimum-cooling-days <integer>]`` - Volume Tiering Minimum
Cooling Days [.advaced-priv]#(privilege: advanced)#::
This parameter displays the minimum number of days that user data blocks
of the volume must be cooled before they can be considered cold and tiered
out to the cloud tier. For volumes hosted on FabricPools, this parameter
is used for tiering purposes and does not affect the reporting of inactive
data. For volumes hosted on non-FabricPools, this parameter affects the
cooling window used for reporting inactive data. The default minimum
cooling period for the ``_snapshot-only_`` policy is 2 days and for the
``_auto_`` policy is 31 days.

``[-performance-tier-inactive-user-data {<integer>[KB|MB|GB|TB|PB]}]`` -
Performance Tier Inactive User Data::
If this parameter is specified, the command displays the amount of
inactive user data stored in the performance tier that could be tiered out
to a cloud tier if the volume is in a FabricPool and for which the
``_auto_`` tiering policy has been specified. For more information see the
tiering-policy parameter.

``[-performance-tier-inactive-user-data-percent <percent>]`` - Performance
Tier Inactive User Data Percent::
If this parameter is specified, the command displays the percentage of
inactive user data in the performance tier.

``[-tiering-object-tags <text>,...]`` - Tags to be Associated with Objects
Stored on a FabricPool::
This optional parameter specifies the tiering object tags to be associated
with objects stored on a FabricPool.

``[-needs-object-retagging {true|false}]`` - Does the Object Tagging
Scanner Need to Run on This Volume::
This optional parameter specifies if an object tagging scanner needs to
```

```
run again for the volume.

``[-is-analytics-supported {true|false}]`` - Is File System Analytics
Supported::
If this parameter is specified, the command displays information only
about volumes that are supported (true) or not supported (false). This
field indicates whether or not file system analytics is supported on the
volume. If file system analytics is not supported, the reason will be
specified in the ``_analytics_unsupported_reason_`` field.

``[-analytics-unsupported-reason <text>]`` - Reason File System Analytics
is not Supported::
If this parameter is specified, the command displays information only
about volumes that are not supported with the specified reason. If file
system analytics is not supported on the volume, this field provides the
reason why.

``[-analytics-state <Analytics State>]`` - File System Analytics State::
If this parameter is specified, the command displays information only
about volumes with file system analytics ``_on_`` , ``_off_`` or
``_initializing_`` . If this value is ``_on_`` , ONTAP collects extra file
system analytics information for all directories on the volume. There will
be a slight impact to I/O performance to collect this information. If this
value is ``_off_`` , file system analytics information is not collected
and not available to be viewed. If this value is ``_initializing_`` , that
means file system analytics was recently turned on, and the initialization
scan to gather information for all existing files and directories is
currently running.

``[-analytics-scan-progress <percent>]`` - File System Analytics Scan
Progress::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified file system analytics
progress percentage. This value will only be set if a file system
analytics initialization scan is in progress.

``[-analytics-files-scanned <integer>]`` - File System Analytics Files
Scanned Progress::
If this parameter is specified, the command displays information only
about the volume or volumes that have the number of files scanned in the
file system. This value will only be set if a file system analytics
initialization scan is in progress.

``[-analytics-total-files <integer>]`` - File System Analytics Total
Files::
If this parameter is specified, the command displays information only
```

```
about the volume or volumes that have the total files in the file system
analytics progress percentage. This value will only be set if a file
system analytics initialization scan is in progress.

``[-activity-tracking-state <Activity Tracking State>]`` - Activity
Tracking State::
If this parameter is specified, the command displays information only
about volumes with volume activity tracking ``_on_`` or ``_off_`` . If
this value is ``_on_`` , ONTAP tracks volume activity in real time to
provide a more detailed view. There will be a slight impact to I/O
performance to collect this information. If this value is ``_off_`` ,
activity is not tracked and not available to be viewed.

``[-is-activity-tracking-supported {true|false}]`` - Is Activity Tracking
Supported::
If this parameter is specified, the command displays information only
about volumes that are supported (true) or not supported (false). This
field indicates whether or not volume activity tracking is supported on
the volume. If volume activity tracking is not supported, the reason will
be specified in the ``_activity_tracking_unsupported_reason_`` field.

``[-activity-tracking-unsupported-reason <text>]`` - Reason Activity
Tracking Is Not Supported::
If this parameter is specified, the command displays information only
about volumes that are not supported with the specified reason. If volume
activity tracking is not supported on the volume, this field provides the
reason why.

``[-is-smbc-master {true|false}]`` - Is SnapMirror Active Sync Master::
If this parameter is specified, it displays if the volume is acting as a
SnapMirror active sync master.

``[-is-smbc-failover-capable {true|false}]`` - Is SnapMirror Active Sync
Failover Capable::
This parameter specifies if volume is capable of SnapMirror active sync
failover.

``[-smbc-consensus {Awaiting-consensus|Consensus|No-consensus}]`` -
SnapMirror Active Sync Consensus::
This parameter specifies the SnapMirror active sync consensus value.
+
Consensus indicated if cluster has to serve IO. Possible consensus values:
+
  * Consensus Awaiting: Allow IO
  * Consensus: Allow IO
  * No consensus: Do not allow IO
```

```
``[-anti-ransomware-state {disabled|enabled|dry-run|paused|dry-run-
paused|enable-paused|disable-in-progress}]`` - Anti-ransomware State::
If this parameter is specified, the command displays information only
about the volumes that have the specified Anti-ransomware-state. The
possible values are ``_disabled_`` , ``_enabled_`` , ``_dry-run_`` ,
``_dry-run-paused_`` , ``_enable-paused_`` or ``_disable-in-progress_`` .
The value ``_disabled_`` is used for volumes which do not have Anti-
ransomware feature enabled. The value is ``_enabled_`` when Anti-
ransomware feature is enabled on the volume. If the value is ``_dry-run_``
, the volume is in Anti-ransomware evaluation mode. The value is ``_dry-
run-paused_`` when the Anti-ransomware feature is paused from evaluation
mode on the volume. The value is ``_enable-paused_`` when the Anti-
ransomware feature is paused on the volume. The value is ``_disable-in-
progress_`` when disable work is in progress on the volume.

``[-granular-data {disabled|basic|advanced}]`` - Granular data::
If this parameter is specified, the command displays information only
about volumes that match the specified value for whether data storage on
the volume is granular or not.

``[-atime-update-period <integer>]`` - Access Time Update Period (Seconds)
[.advaced-priv]#(privilege: advanced)#::
"If this parameter is specified, the command displays information only
about volumes that match the specified value for the atime update period."

``[-snapshot-locking-enabled {true|false}]`` - Enable Snapshot Copy
Locking::
"If this parameter is specified, the command displays information only
about volumes that match the specified `-snapshot-locking-enabled` value.
A volume with locked snapshots cannot be deleted until the `-expiry-time`
has passed."

``[-expiry-time <text>]`` - Expiry Time::
 If this parameter is specified, the command displays all the volumes that
match the specified `-expiry-time` value.

``[-compliance-clock-time <text>]`` - ComplianceClock Time::
 If this parameter is specified, the command displays all the volumes that
match the specified `-compliance-clock-time` value.

``[-is-large-size-enabled {true|false}]`` - Are Large Size Volumes and
Files Enabled::
 If this parameter is specified, the command displays all the volumes that
match the specified `-is-large-size-enabled` value.
```

```
``[-is-preserve-unlink-enabled {true|false}]`` - Is Preserve Unlink
Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about volumes that match the specified `-is-preserve-unlink-enabled`
value.

``[-is-cloud-write-enabled {true|false}]`` - Is Cloud Write Enabled
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays all the volumes that
match the specified `-is-cloud-write-enabled` value.

``[-aggressive-readahead-mode
{none|file_prefetch|sequential_read|cross_file_sequential_read}]`` -
Aggressive readahead mode [.advaced-priv]#(privilege: advanced)#::
 If this parameter is specified, the command displays all the volumes that
match the specified `-aggressive-readahead-mode` value.

``[-in-consistency-group {true|false}]`` - If this Volume is part of a
Consistency Group::
This parameter specifies if this volume is a associated with a consistency
group.

``[-total-metadata {<integer>[KB|MB|GB|TB|PB]}]`` - Total Metadata Size::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified amount of space used
by metadata in the volume. This field will not include the metadata of the
volume which is residing inside the aggregate.

``[-total-metadata-footprint {<integer>[KB|MB|GB|TB|PB]}]`` - Total
Metadata Footprint Used in Aggregate::
If this parameter is specified, the command displays information only
about volume or volumes that have the specified amount of space used by
volume metadata footprint in the aggregate. The value of this field will
be higher than the field 'total-metadata' from the
xref:{relative_path}volume-show-space.html[volume show-space] command
output since this includes additional metadata of the volume in the
aggregate.

== Examples

The following example displays information about all volumes on the
Vserver named vs1:
```

```
cluster1::> volume show -vserver vs1
Vserver Volume Aggregate State Type Size Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
```

vs1 vol1 aggr1 online RW 2GB 1.9GB 5%
vs1 vol1_dr aggr0_dp online DP 200GB 160.0GB 20%
vs1 vol2 aggr0 online RW 150GB 110.3GB 26%
vs1 vol2_dr aggr0_dp online DP 150GB 110.3GB 26%
vs1 vol3 aggr1 online RW 150GB 120.0GB 20%
vs1 vol3_dr aggr1_dp online DP 150GB 120.0GB 20%
vs1 vol4 aggr1 online RW 200GB 159.8GB 20%
vs1 vol4_dr aggr1_dp online DP 200GB 159.8GB 20%
vs1 vol5 aggr2 online RW 200GB 102.3GB 48%
vs1 vol5_dr aggr2_dp online DP 200GB 102.3GB 48%
vs1 vol6 aggr2 online RW 150GB 117.2GB 21%
vs1 vol6_dr aggr2_dp online DP 150GB 117.2GB 21%
vs1 vol7 aggr3 online RW 150GB 118.5GB 20%
vs1 vol7_dr aggr3_dp online DP 150GB 118.5GB 20%
vs1 vol8 aggr3 online RW 150GB 90.03GB 39%
vs1 vol8_dr aggr3_dp online DP 150GB 90.03GB 39%
vs1 vol9 aggr4 online RW 150GB 43.67GB 70%
vs1 vol9_dr aggr4_dp online DP 150GB 43.67GB 70%
vs1 vol10 aggr4 online RW 150GB 108.7GB 27%
vs1 vol10_dr aggr4_dp online DP 150GB 108.7GB 27%
vs1 vol11 aggr5 online RW 250GB 45.65GB 81%
vs1 vol11_dr aggr5_dp online DP 250GB 45.65GB 81%
22 entries were displayed.

```
The following example displays detailed information about a volume named
vol1 on an SVM named vs1:
```

cluster1::*> volume show -vserver vs1 -volume vol1
Vserver Name: vs1
Volume Name: vol1
Aggregate Name: aggr1
Volume Size: 30MB
Volume Data Set ID: 1026
Volume Master Data Set ID: 2147484674
Volume State: online
Volume Type: RW
Volume Style: flex
Is Cluster Volume: true
Is Constituent Volume: false
Export Policy: default
User ID: root
Group ID: daemon
Security Style: mixed
Unix Permissions: ---rwx------
Junction Path: -
Junction Path Source: -
Junction Active: -
Junction Parent Volume: -
Comment:
Available Size: 23.20MB
Filesystem Size: 30MB
Total User-Visible Size: 28.50MB

Used Size: 5.30MB
Used Percentage: 22%
Volume Nearly Full Threshold Percent: 95%
Volume Full Threshold Percent: 98%
Maximum Autosize (for Flexvol volumes only): 8.40GB
Minimum Autosize: 30MB
Autosize Grow Threshold Percentage: 85%
Autosize Shrink Threshold Percentage: 50%
Autosize Mode: off
Autosize Enabled (for Flexvol volumes only): false
Total Files (for user-visible data): 217894
Files Used (for user-visible data): 98
Space Guarantee Style: volume
Space Guarantee In Effect: true
Snapshot Directory Access Enabled: true
Space Reserved for Snapshot Copies: 5%
Snapshot Reserve Used: 98%
Snapshot Policy: default
Creation Time: Mon Jul 08 10:54:32 2013
Language: C.UTF-8
Clone Volume: false
Node name: cluster-1-01
NVFAIL Option: off
Force NVFAIL on MetroCluster Switchover: off
Is File System Size Fixed: false
Extent Option: off
Reserved Space for Overwrites: 0B
Fractional Reserve: 100%
Primary Space Management Strategy: volume_grow
Read Reallocation Option: space-optimized
Inconsistency in the File System: false
Is Volume Quiesced (On-Disk): false
Is Volume Quiesced (In-Memory): false
Transition Operation State: none
Copied for Transition: false
Transitioned: true
Volume Contains Shared or Compressed Data: false
Efficiency Policy: default
UUID of the Efficiency Policy: b0f36cd7-e7bc-11e2-9994-123478563412
Space Saved by Storage Efficiency: 0B
Percentage Saved by Storage Efficiency: 0%
Space Saved by Deduplication: 0B
Percentage Saved by Deduplication: 0%
Space Shared by Deduplication: 0B
Space Saved by Compression: 0B
Percentage Space Saved by Compression: 0%
Volume Size Used by Snapshot Copies: 1.48MB
Block Type: 64-bit
Is Volume Moving: false
Flash Pool Caching Eligibility: read-write
Flash Pool Write Caching Ineligibility Reason: -
Managed By Storage Service: -
Enable Object Store: -
Constituent Volume Role: -

Is cft precommit: false
QoS Policy Group Name: -
Caching Policy Name: auto
Is Volume Move in Cutover Phase: false
Number of Snapshot Copies in the Volume: 10
VBN_BAD may be present in the active filesystem: false
Is Eligible for Auto Balance Aggregate: -
Is Volume on a hybrid aggregate: false
Total Physical Used Size: 4.55MB
Physical Used Percentage: 14%
FlexGroup volume Master Data Set ID: -
FlexGroup volume Index: -
UUID of the FlexGroup volume: -
Maximum size of a FlexGroup volume Constituent: -
Inofile Version: 3
List of Nodes: -
Is Volume FlexGroup: false
SnapLock Type: -
Vserver DR Protection: -
Healthy: true
Unhealthy Reason:
Is Fenced for Protocol Access: false
Protocol Access Fence Owner: -
Is SnapMirror active sync Master: false
Is SnapMirror active sync Failover Capable: false
SnapMirror active sync Consensus: -

```
== Related Links

* xref:{relative_path}volume-show-space.html[volume show-space]




[[ID78eb2582f19d45e3cdf1c778ecafed6a]]
= volume size
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set/Display the size of the volume.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume size` command allows the user to set or display the volume
```

```
size. If ``_new-size_`` is not specified then the current volume size is
displayed.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter can be used to specify the Vserver on which the volume is
located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the volume for which the user wants to set or
display the size.

``[-new-size {<integer>[KB|MB|GB|TB|PB]}]`` - [+|-]<New Size>::
This optional parameter specifies the size of the volume. It can be used
to set the volume size to a particular number or grow/shrink the size by a
particular amount. The size is specified as a number (preceded with a sign
for relative growth/shrinkage) followed by a unit designation: k
(kilobytes), m (megabytes), g (gigabytes), or t (terabytes). If the unit
designation is not specified, bytes are used as the unit, and the
specified number is rounded up to the nearest 4 KB. The minimum size for a
flexible volume is 20 MB, and the maximum size depends on hardware
platform and free space in the containing aggregate. If the volume's space
guarantee is currently disabled, its size cannot be increased.

== Examples

The following example shows the size of a volume called vol1.
```

cluster1::> vol size vol1
(volume size)
vol size: Flexible volume 'vs1:vol1' has size 2g.

```
The following example sets the size of a volume called vol1 to 1GB.
```

cluster1::> vol size vol1 1g
(volume size)
vol size: Flexible volume 'vs1:vol1' size set to 1g.

```
The following example increases the size of a volume called vol1 by 500MB.
```

cluster1::> vol size vol1 +500m
(volume size)
vol size: Flexible volume 'vs1:vol1' size set to 1.49g.

```
The following example decreases the size of a volume called vol1 by 250MB.
```

cluster1::> vol size vol1 -250m
(volume size)
vol size: Flexible volume 'vs1:vol1' size set to 1.24g.

```
[[ID18a6d24407e8c1f316e173870b0375e7]]
= volume transition-prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Verifies that there are no volumes actively transitioning from 7-mode to
clustered Data ONTAP, and configures the transition feature for downgrade.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume transition-prepare-to-downgrade` command is used to verify
that a volume is not currently being transitioned from 7-Mode to ONTAP.
This check must be done before reverting or downgrading a node.

== Parameters

== Examples

The following example executes the `volume transition-inactive-verify`
command during active volume transitions.

cluster1::*> volume transition-prepare-to-downgrade

Error: command failed: Downgrade cannot proceed because one or more
volumes are transitioning from 7-Mode to ONTAP.

You must complete or cancel all transitions. View the transitioning
volumes that are preventing the downgrade by using the (privilege:
advanced) command "lun transition show -status active -fields vserver,
volume".

To complete a transition, break the SnapMirror transition relationship
using the command "snapmirror break -destination-path <destination-path>".
```

```
For detailed information about transitioning, refer to the "7-Mode Data
Transition Using SnapMirror" guide.

To cancel a transition, contact technical support.




[[ID51d9d37397df4a7760bfcaea340e2769]]
= volume unmount
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Unmount a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume unmount` command unmounts a volume from its parent volume. The
volume can be remounted at the same or a different location by using the
xref:{relative_path}volume-mount.html[volume mount] command.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the volume that is to be unmounted.

== Examples

The following example unmounts a volume named vol2 on a Vserver named vs0:
```

node::> volume unmount -vserver vs0 -volume vol2

```
== Related Links

* xref:{relative_path}volume-mount.html[volume mount]
```

```
= volume activity-tracking commands

:leveloffset: +1


[[ID31e67c7c222e04ad96229000d1370de7]]
= volume activity-tracking off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable activity tracking for a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume activity-tracking off` command disables volume activity
tracking on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which volume activity tracking is
being disabled.

``[-foreground <true>]`` - Run Operation in the Foreground::
If true, the command doesn't complete until the volume activity tracking
`state` is ``_off_`` . If false, the command returns immediately, and the
volume activity tracking `state` can be tracked using the
xref:{relative_path}volume-activity-tracking-show.html[volume activity-
tracking show] command.

== Examples

The following examples disable volume activity tracking on a volume.
```

cluster::*> volume activity-tracking off -vserver vs1 -volume fv1

```
```

cluster::*> volume activity-tracking off -vserver vs1 -volume fg1

```
== Related Links

* xref:{relative_path}volume-activity-tracking-show.html[volume activity-
tracking show]




[[IDf442ffb7ec4433390c1ae892b0af06ef]]
= volume activity-tracking on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable activity tracking for a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume activity-tracking on` command enables volume activity tracking
on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which volume activity tracking is
being enabled.

``[-foreground <true>]`` - Run Operation in the Foreground::
If true, the command doesn't complete until the volume activity tracking
`state` is ``_on_`` . If false, the command returns immediately, and the
volume activity tracking `state` can be tracked using the
xref:{relative_path}volume-activity-tracking-show.html[volume activity-
tracking show] command.

== Examples

The following examples enable volume activity tracking on a volume.
```

cluster::*> volume activity-tracking on -vserver vs1 -volume fv1

```
```

cluster::*> volume activity-tracking on -vserver vs1 -volume fg1

```
== Related Links

* xref:{relative_path}volume-activity-tracking-show.html[volume activity-
tracking show]




[[ID1a4a1f8f8eb0c35832f74884c60f2607]]
= volume activity-tracking show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display activity tracking information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume activity-tracking show` command displays the volume activity
tracking information for a volume. By default, the command displays the
following information:

* Vserver name
* Volume name
* Activity tracking state



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
If this parameter and the `-volume` parameter are specified, the command
displays volume activity tracking information about the specified volume.
If this parameter is specified by itself, the command displays information
about volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter and the `-vserver` parameter are specified, the command
displays volume activity tracking information about the specified volume.
If this parameter is specified by itself, the command displays information
about all volumes matching the specified name.

``[-state <Activity Tracking State>]`` - State::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified volume activity
tracking state.

== Examples

The following examples display the volume activity tracking information
for a volume.
```

cluster::*> volume activity-tracking show -vserver vs1 -volume fv1
Vserver Name: vs1
Volume Name: fv1
State: on

```

```

cluster::*> volume activity-tracking show -vserver vs1 -volume fg1
Vserver Name: vs1
Volume Name: fg1
State: on

```
:leveloffset: -1


= volume analytics commands

:leveloffset: +1


[[ID0407a4494c794656963c31d2e3772af5]]
= volume analytics off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable collection of file system analytics for a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume analytics off` command disables file system analytics on a
volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which file system analytics is
being disabled.

``[-foreground <true>]`` - Run Operation in the Foreground::
If true, the command doesn't complete until the `analytics-state` of the
volume is ``_off_`` . If false, the command returns immediately, and the
`analytics-state` can be tracked using the xref:{relative_path}volume-
analytics-show.html[volume analytics show] command.

== Examples

The following example disables file system analytics on a flexible volume
named "flexvol" in Vserver "vs1":
```

```
cluster::*> volume analytics off -vserver vs1 -volume flexvol
```

```
== Related Links

* xref:{relative_path}volume-analytics-show.html[volume analytics show]




[[ID81bda1e1784dcc086bcf9add942082e]]
= volume analytics on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable collection of file system analytics for a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume analytics on` command enables file system usage analytics on a
volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which file system usage analytics
is being enabled.

``[-foreground <true>]`` - Run Operation in the Foreground::
If true, the command doesn't complete until the `analytics-state` of the
volume is ``_on_`` . If false, the command returns immediately, and the
`analytics-state` and `scan-progress` can be tracked using the
xref:{relative_path}volume-analytics-show.html[volume analytics show]
command.

== Examples

The following example enables file system usage analytics on a flexible
volume named "flexvol" in Vserver "vs1":
```

cluster::*> volume analytics on -vserver vs1 -volume flexvol

```
== Related Links

* xref:{relative_path}volume-analytics-show.html[volume analytics show]




[[ID20b9466552954dc978b2c7abe660a546]]
= volume analytics show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display file system analytics information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume analytics show` command displays the file system analytics
information for a volume. By default, the command displays the following
information:

* Vserver name
* Volume name
* File system analytics state



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
```

```
If this parameter and the `-volume` parameter are specified, the command
displays file system analytics information about the specified volume. If
this parameter is specified by itself, the command displays information
about volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter and the `-vserver` parameter are specified, the command
displays file system analytics information about the specified volume. If
this parameter is specified by itself, the command displays information
about all volumes matching the specified name.

``[-state <Analytics State>]`` - State::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified file system analytics
state.

``[-scan-progress <percent>]`` - Scan Progress::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified file system analytics
progress percentage. This value will only be set if a file system
analytics initialization scan is in progress.

``[-tracked-dirs <integer>]`` - Tracked Directories::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of tracked
directories.

``[-tracked-files <integer>]`` - Tracked Files::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of tracked
files.

``[-tracked-items <integer>]`` - Tracked Items::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of tracked
items. An item is either a file or a directory.

``[-top-level-dirs <integer>]`` - Top Level Directories::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of top level
directories.

``[-average-files-per-dir <integer>]`` - Average Files Per Directory::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified average number of
files per directory.
```

```
``[-files-scanned <integer>]`` - Number of Files Scanned::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified number of files
scanned. This value will only be set if a file system analytics
initialization scan is in progress.

``[-total-files <integer>]`` - Total Number of Files::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified total number of files.
This value will only be set if a file system analytics initialization scan
is in progress.

== Examples

The following example displays the file system analytics information on a
flexible volume named "flexvol" in Vserver "vs1":
```

cluster::*> volume analytics show -vserver vs1 -volume flexvol
Vserver Name: vs1
Volume Name: flexvol
State: on
Scan Progress: -

```
[[ID86d6bdbda0927464fa1f0982ef04ee6d]]
= volume analytics initialization pause
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Pause the analytics file system scan

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume analytics initialization pause` command pauses the file system
analytics scan.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which file system analytics
initialization is being paused.

== Examples

The following example pauses the file system analytics scan on a volume
named "vol1" in Vserver "vs1":
```

cluster::*> volume analytics initialization pause -vserver vs1 -volume vol1

```
[[ID1e50d21931440e7f9e954fd8d7bcdfba]]
= volume analytics initialization resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume the analytics file system scan

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume analytics initialization resume` command resumes the file
system analytics scan.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which file system analytics
initialization is being resumed.

== Examples

The following example resumes the file system analytics scan on a volume
named "vol1" in Vserver "vs1":
```

cluster::*> volume analytics initialization resume -vserver vs1 -volume vol1

```
[[ID36cc8ac044396a2df531eb470341fdbc]]
= volume analytics initialization show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display file system analytics information

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

The `volume analytics initialization show` command displays the file
system analytics initialization information for a volume. By default, the
command displays the following information:

* Vserver name
* Volume name
* File system analytics initialization state



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
If this parameter is specified, the command displays information only
about the volumes in vservers that match the specified value for `-
vserver`

``[-volume <volume name>]`` - Volume Name::
If this parameter is specified, the command displays information only
about the volumes that match the specified value for `-volume`

``[-state <Analytics Initialization State>]`` - State::
If this parameter is specified, the command displays information only
about the volume or volumes that have the specified file system analytics
initialization state.

== Examples

The following example displays the file system analytics initialization
information on a volume named "vol1" in Vserver "vs1":
```

cluster::*> volume analytics initialization show -vserver vs1 -volume vol1
Vserver Name: vs1

Volume Name: vol1
State: Running

```
:leveloffset: -1


= volume clone commands

:leveloffset: +1


[[ID3f09a81966c690110094a2b0e639f698]]
= volume clone create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a FlexClone volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume clone create` command creates a FlexClone volume on the
aggregate containing the specified parent volume. This command is
supported for flexible volumes or FlexGroup volumes. The maximum volume
clone hierarchy depth is 500 and the default depth is 60. You can
optionally specify the following attributes for the new FlexClone volume:

* Vserver on which the parent volume resides
* Name of the FlexClone parent snapshot
* Junction path where FlexClone volume should be mounted
* State of the junction path
* Space guarantee style (none, volume or file)
* Comment
* Whether the `volume clone create` command runs as a foreground or
background process
* Key Manager Attribute



== Parameters
```

```
``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the FlexClone volume is to
be created. If only one data Vserver exists, you do not need to specify
this parameter.

``-flexclone <volume name>`` - FlexClone Volume::
This parameter specifies the name of the FlexClone volume. The name must
be unique within the hosting Vserver.

``[-type {RW|DP}]`` - FlexClone Type::
This parameter specifies the type of FlexClone volume. A read-only
FlexClone volume is created if you specify the ``_type_`` as DP; otherwise
a read-write FlexClone volume is created.

``[-parent-vserver <vserver name>]`` - FlexClone Parent Vserver::
This parameter specifies the name of the Vserver to which the FlexClone
parent volume belongs. If it is different from the Vserver on which the
FlexClone volume is to be created, then the FlexClone volume inherits the
export policies from the residing Vserver, and not from the FlexClone
parent volume.

``-b, -parent-volume <volume name>`` - FlexClone Parent Volume::
This parameter specifies the name of parent volume from which the
FlexClone clone volume is derived.

``[-parent-snapshot <snapshot name>]`` - FlexClone Parent Snapshot::
This specifies the name of the parent snapshot from which the FlexClone
clone volume is derived.

``[-junction-path <junction path>]`` - Junction Path::
This specifies the junction path at which the new FlexClone clone volume
should be mounted.

``[-junction-active {true|false}]`` - Junction Active::
This optionally specifies whether the volume's junction path is active.
The default setting is `true` . If the junction path is inactive, the
volume does not appear in the Vserver's namespace. This parameter is
available only at the advanced privilege level and higher.

``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
This optionally specifies the space guarantee style for the FlexClone
volume. A value of ``_volume_`` reserves space on the aggregate for the
entire volume. A value of ``_none_`` reserves no space on the aggregate,
meaning that writes can fail if the aggregate runs out of space. The
default setting is inherited from the parent volume.
```

```
``[-comment <text>]`` - Comment::
This optionally specifies a comment for the FlexClone volume.

``[-foreground {true|false}]`` - Foreground Process::
This optionally specifies whether the FlexClone volume create operation
runs as a foreground process. The default setting is `true` (that is, the
operation runs in the foreground).

{ ``[-qos-policy-group <text>]`` - QoS Policy Group Name::
This parameter optionally specifies which QoS policy group to apply to the
FlexClone volume. The policy group defines measurable service level
objectives (SLOs) that apply to the storage objects with which the policy
group is associated. If you do not assign a policy group to the FlexClone
volume, the system does not monitor and control the traffic to the volume.

| ``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name
}::
This optionally specifies which QoS adaptive policy group to apply to the
volume. This policy group defines measurable service level objectives
(SLOs) and Service Level Agreements (SLAs) that adjust based on the volume
allocated space or used space. This parameter is not supported on
FlexGroup volumes.

``[-caching-policy <text>]`` - Caching Policy Name::
 This optionally specifies the caching policy to apply to the volume. A
caching policy defines how the system caches this volume's data in Flash
Cache modules. If a caching policy is not assigned to this volume, the
system uses the caching policy that is assigned to the containing Vserver.
If a caching policy is not assigned to the containing Vserver, the system
uses the default cluster-wide policy. The available caching policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read, and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.
  * noread-random_write - Write caches all randomly overwritten user data
```

```
blocks. It does not do any read caching.
  * meta-random_write - Read caches all metadata and write caches randomly
overwritten user data blocks.
  * random_read_write-random_write - Read caches all metadata, randomly
read and randomly written user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read-random_write - Read caches all metadata, randomly read, and
sequentially read user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read_random_write-random_write - Read caches all metadata,
randomly read, sequentially read and randomly written user data. It also
write caches randomly overwritten user data blocks.
  * all-random_write - Read caches all data blocks read and written. It
also write caches randomly overwritten user data blocks.

+
Note that in a caching-policy name, a hyphen (-) separates read and write
policies. Default caching-policy is auto.

``[-vserver-dr-protection {protected|unprotected}]`` - Vserver DR
Protection::
 This optionally specifies whether the volume should be protected by
Vserver level SnapMirror. This parameter is applicable only if the Vserver
is the source of a Vserver level SnapMirror relationship. By default the
clone volume will inherit this value from the parent volume.

``[-uid <integer>]`` - Volume-Level UID::
This parameter optionally specifies a volume-level user ID (UID). All
files and directories in a FlexClone volume will inherit this UID.

``[-gid <integer>]`` - Volume-Level GID::
This parameter optionally specifies a volume-level group ID (GID). All
files and directories in a FlexClone volume will inherit this GID.

``[-snaplock-type {non-snaplock|compliance|enterprise}]`` - SnapLock
Type::
This parameter optionally specifies the SnapLock type for the FlexClone
volume. The `snaplock-type` can be `compliance` , `enterprise` or `non-
snaplock` . By default, the `snaplock-type` will be same as the parent
volume.

``[-key-manager-attribute <text>]`` - Key Manager Attribute::
This optional parameter specifies an additional key manager attribute that
should be an identifier-value pair separated by '=', ex. CRN=my-unique-
value
+
```

```
The following identifiers are currently supported:
+
  * CRN




== Examples

The following command creates a FlexClone volume ``_fc_vol_1_`` from
parent volume ``_fv2_`` on Vserver ``_vs1_`` and the job runs as a
foreground process.
```

cluster1::> volume clone create -vserver vs1 -flexclone fc_vol_1 -parent-volume fv2 -junction-active true -foreground true -comment "Testing FlexClone creation"
cluster1::> volume clone show fc_vol_1 -vserver vs1
Vserver Name: vs1
FlexClone Volume: fc_vol_1
FlexClone Parent Volume: fv2
FlexClone Parent Snapshot: clone_fc_vol_1.0
Junction Path: -
Junction Active: -
Space Guarantee Style: volume
Space Guarantee In Effect: true
FlexClone Aggregate: test_aggr
FlexClone Data Set ID: 1046
FlexClone Master Data Set ID: 2147484694
FlexClone Size: 19MB
Used Size: 108KB
Split Estimate: 0.00B
Inodes processed: -
Total Inodes: -
Percentage complete: -
Blocks Scanned: -
Blocks Updated: -
Comment: Testing FlexClone creation

```
[[ID38e8f07e4424972866e979d274dc18ac]]
= volume clone show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of FlexClones

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `volume clone show` command displays information about FlexClone clone
volumes. This command is only supported for flexible volumes. By default,
the command displays the following information about all FlexClone volume
clones:

* Vserver name
* FlexClone volume name
* Parent volume name
* Parent snapshot name
* Whether a FlexClone volume is online or offline

To display detailed information about all FlexClone volumes, run the
command with the `-instance` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields to be displayed.

| ``[-estimate  ]``::
Displays an estimate of the free disk space required in the aggregate to
split the indicated clone volume from its underlying parent volume. The
value reported may differ from the space actually required to perform the
split, especially if the clone volume is changing when the split is being
performed.

| ``[-instance  ]`` }::
Displays detailed information about FlexClone volumes. If `-flexclone` is
also specified, the command displays detailed information about the
FlexClone volume.

``[-vserver <vserver name>]`` - Vserver Name::
Selects summary information for the FlexClone volumes on the specified
Vserver. If `-flexclone` is also specified, the command displays detailed
information about the specified FlexClone volume.

``[-flexclone <volume name>]`` - FlexClone Volume::
Selects summary information for the specified FlexClone volume. If `-
vserver` is also specified, the command displays detailed information
about the specified FlexClone volume.

``[-type {RW|DP}]`` - FlexClone Type::
Selects information for the specified type of FlexClone volume. The type
can be specified as either read-only (DP) or read-write (RW).
```

```
``[-parent-vserver <vserver name>]`` - FlexClone Parent Vserver::
Selects summary information for the FlexClone volumes that are clone
volumes in the specified parent Vserver.

``[-b, -parent-volume <volume name>]`` - FlexClone Parent Volume::
Selects summary information for the FlexClone volumes that are clones of
the specified parent volume.

``[-parent-snapshot <snapshot name>]`` - FlexClone Parent Snapshot::
Selects summary information for the FlexClone volumes that are clones of
the parent volume to which the specified snapshot belongs.

``[-state {online|restricted|offline|force-online|force-offline|mixed}]``
- FlexClone Volume State::
Selects summary information for the FlexClone volumes that are in the
specified state.

``[-junction-path <junction path>]`` - Junction Path::
Selects summary information for the FlexClone volumes that have the
specified junction path.

``[-junction-active {true|false}]`` - Junction Active::
Selects summary information for the FlexClone volumes that have the
specified junction path status.

``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
If this parameter is specified, the command displays information only
about the volumes that have the specified space guarantee style.

``[-space-guarantee-enabled {true|false}]`` - Space Guarantee In Effect::
Selects summary information for the FlexClone volumes that have the
specified space-guarantee setting.

``[-aggregate <aggregate name>]`` - FlexClone Aggregate::
Selects summary information for the FlexClone volumes that reside on the
specified storage aggregate.

``[-dsid <integer>]`` - FlexClone Data Set ID::
Selects summary information for the FlexClone volumes that have the
specified Data Set ID.

``[-msid <integer>]`` - FlexClone Master Data Set ID::
Selects summary information for the FlexClone volumes that have the
specified Master Data Set ID.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - FlexClone Size::
Selects summary information for the FlexClone volumes that have the
```

```
specified size.

``[-used {<integer>[KB|MB|GB|TB|PB]}]`` - Used Size::
Selects summary information for the FlexClone volumes that have the
specified amount of used space.

``[-split-estimate {<integer>[KB|MB|GB|TB|PB]}]`` - Split Estimate::
Selects summary information for the FlexClone volumes that require the
specified amount of free disk space for splitting from the parent.

``[-blocks-scanned <integer>]`` - Blocks Scanned::
 Selects summary information for the FlexClone volumes that have the
specified number of blocks scanned for splitting the FlexClone volume from
its parent volume.

``[-blocks-updated <integer>]`` - Blocks Updated::
Selects summary information for the FlexClone volumes that have the
specified number of blocks updated for after splitting the FlexClone
volume from its parent volume.

``[-comment <text>]`` - Comment::
Selects summary information for the FlexClone volumes that have the
specified comment.

``[-qos-policy-group <text>]`` - QoS Policy Group Name::
Selects summary information for the FlexClone volumes that have the
specified QoS policy group.

``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name::
Selects summary information for the FlexClone volumes that have the
specified QoS adaptive policy group.

``[-caching-policy <text>]`` - Caching Policy Name::
 This optionally specifies the caching policy to apply to the volume. A
caching policy defines how the system caches this volume's data in Flash
Cache modules. If a caching policy is not assigned to this volume, the
system uses the caching policy that is assigned to the containing Vserver.
If a caching policy is not assigned to the containing Vserver, the system
uses the default cluster-wide policy. The available caching policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
```

```
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read, and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.
  * noread-random_write - Write caches all randomly overwritten user data
blocks. It does not do any read caching.
  * meta-random_write - Read caches all metadata and write caches randomly
overwritten user data blocks.
  * random_read_write-random_write - Read caches all metadata, randomly
read and randomly written user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read-random_write - Read caches all metadata, randomly read, and
sequentially read user data blocks. It also write caches randomly
overwritten user data blocks.
  * all_read_random_write-random_write - Read caches all metadata,
randomly read, sequentially read and randomly written user data. It also
write caches randomly overwritten user data blocks.
  * all-random_write - Read caches all data blocks read and written. It
also write caches randomly overwritten user data blocks.

+
Note that in a caching-policy name, a hyphen (-) separates read and write
policies. Default caching-policy is auto.

``[-parent-vol-type <volAccessType>]`` - Parent volume type [.advaced-
priv]#(privilege: advanced)#::
Selects summary information for the FlexClone volumes that are clones of
the parent volumes with the specified type.

``[-flexclone-used-percent <percent>]`` - FlexClone Used Percentage::
Selects summary information for the FlexClone volumes that have the
specified percentage of used space.

``[-vserver-dr-protection {protected|unprotected}]`` - Vserver DR
Protection::
Selects summary information for the FlexClone volumes that have the
specified type of Vserver SnapMirror protection. This parameter is
applicable only if the Vserver is the source of a Vserver level SnapMirror
relationship.

``[-block-percentage-complete <integer>]`` - Percentage Complete::
Selects summary information for the FlexClone volumes that have specified
```

```
percentage of Blocks processed for splitting the FlexClone volume from its
parent volume.

``[-uid <integer>]`` - Volume-Level UID::
Selects summary information for the FlexClone volumes that are created
with the specified volume-level UID.

``[-gid <integer>]`` - Volume-Level GID::
Selects summary information for the FlexClone volumes that are created
with the specified volume-level GID.

``[-flexgroup-uuid <UUID>]`` - UUID of the FlexGroup::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volume constituents that have the
specified FlexGroup volume UUID. This parameter applies to FlexGroup
volumes and FlexGroup volume constituents only.

``[-flexgroup-msid <integer>]`` - FlexGroup Master Data Set ID::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volume constituents that have the
specified FlexGroup volume master data-set ID. This parameter applies to
FlexGroup volumes and FlexGroup volume constituents only.

``[-flexgroup-index <integer>]`` - FlexGroup Index::
If this parameter is specified, the command displays information only
about the FlexGroup volume constituents that have the specified FlexGroup
volume index. This parameter applies to FlexGroup volume constituents
only.

``[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum size of
a FlexGroup Constituent::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volumes that have the specified
maximum constituent size. This parameter applies to FlexGroup volumes
only.

``[-constituent-role <Constituent Roles>]`` - Constituent Volume Role::
If this parameter is specified, the command displays information only
about the constituent volume or volumes that are of the specified role.
This parameter applies to FlexGroup volumes only.

``[-is-flexgroup-constituent-active {true|false}]`` - Is Active FlexGroup
Constituent::
If this parameter is specified, the command displays information only
about the volume or volumes that are either active constituent of the
FlexGroup or not, depending on the value provided.
```

```
``[-is-constituent {true|false}]`` - Is Constituent Volume::
If this parameter is specified, the command displays information only
about volumes that either are or are not constituents of a FlexGroup,
depending on the value provided.

``[-is-flexgroup {true|false}]`` - Is Volume a FlexGroup::
If this parameter is specified, the command displays information only
about the volume or volumes that are either FlexGroup volumes or not,
depending on the value provided.

``[-volume-style-extended {flexvol|flexgroup|flexgroup-constituent}]`` -
Extended Volume Style::
If this parameter is specified, the command displays information only
about the volumes that are configured with the specified extended style.
Possible values are ``_flexvol_`` for FlexVol volumes, ``_flexgroup_`` for
FlexGroup volumes and ``_flexgroup-constituent_`` for FlexGroup
constituents.

``[-snaplock-type {non-snaplock|compliance|enterprise}]`` - SnapLock
Type::
 Selects summary information for the FlexClone volumes that are created
with the specified SnapLock type.

== Examples

The following example displays detailed information about all FlexClone
volumes on Vserver ``_vs0_`` :
```

cluster1::> volume clone show -vserver vs0
(volume clone show)
Vserver FlexClone Parent-Volume Parent-Snapshot
--------- ---------- --------- ---------
vs0 fc_vol_1 test_vol clone_fc_vol_1.0
fc_vol_2 test_vol2 clone_fc_vol_2.0
fc_vol_3 tv9 clone_fc_vol_3.0
tv8 tv7 clone_tv8.0
tv9 test_vol2 clone_tv9.0
5 entries were displayed.

```
The following example displays detailed information about FlexClone volume
``_fc_vol_2_`` on Vserver ``_vs0_`` :
```

cluster1::> volume clone show -vserver vs0 -flexclone fc_vol_2
Vserver Name: vs0
FlexClone Volume: fc_vol_2
FlexClone Parent Volume: test_vol2
FlexClone Parent Snapshot: clone_fc_vol_2.0

Junction Path: -
Junction Active: -
Space Guarantee Style: volume
Space Guarantee In Effect: true
FlexClone Aggregate: test_aggr
FlexClone Data Set ID: 1038
FlexClone Master Data Set ID: 2147484686
FlexClone Size: 47.50MB
Used Size: 128KB
Split Estimate: 0.00B
Inodes processed: -
Total Inodes: -
Inode Percentage complete: -
Blocks Scanned: -
Blocks Updated: -
Comment:
Qos Policy Group Name: pg1
FlexClone Parent Volume Type: RW
Block Percentage complete: -

```
The following example displays summary information about all FlexClone
volumes residing on Vserver ``_vs0_`` along with the fields `msid` ,
`dsid` , `state` and `parent-volume` .
```

cluster1::> volume clone show -vserver vs0 -fields msid, dsid, state, parent-volume
vserver flexclone parent-volume state dsid msid
------- --------- ------------- ------ ---- ----------
vs0 fc_vol_1 test_vol online 1037 2147484685
vs0 fc_vol_3 tv9 online 1039 2147484687
vs0 flex_clone1
fc_vol_1 online 1041 2147484689
vs0 fv_2 fc_vol_1 online 1043 2147484691
vs0 tv9 test_vol2 online 1036 2147484684
5 entries were displayed.

```
The following example displays summary information about all FlexClone
volumes residing on Vserver ``_vs0_`` along with `space-guarantee-enabled`
and `space-guarantee` information about each FlexClone volume.
```

cluster1::> vol clone show -vserver vs0 -fields space-guarantee-enabled, space-guarantee
(volume clone show)
vserver flexclone space-guarantee space-guarantee-enabled
------- --------- --------------- -----------------------
vs0 fc_vol_1 volume true
vs0 fc_vol_3 volume true
vs0 flex_clone1
volume true
vs0 fv_2 volume true
vs0 tv9 volume true
5 entries were displayed.

```
[[ID72fdb6aaa6c0b395309c39096392de84]]
= volume clone sharing-by-split show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the split flexclone volumes with shared physical blocks

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume clone sharing-by-split show` command displays the split
volumes with shared physical blocks. This command is only supported for
flexible volumes. By default, this command displays the following
information:

* Node Name
* Vserver Name
* Volume Name
* Aggregate Name
* Volume State



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
This parameter selects information about the split volumes with shared
physical blocks on this node.

``[-vserver <Vserver Name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
```

```
This parameter selects information about the split volumes with shared
physical blocks on this Vserver.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This parameter selects information about shared physical blocks for this
volume.

``[-aggregate <aggregate name>]`` - Aggregate Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the aggregate associated with the given volume.

== Examples

The following example displays the split volumes with shared physical
blocks in the node:
```

cluster1::> volume clone sharing-by-split show -node node1
Node Vserver Volume Aggregate
--- --- --- ---
node1 vs1 vol_clone1 aggr1

```
The following example displays information about volume ``_vol_clone1_``
residing on vserver ``_vs1_`` :
```

cluster1::> volume clone sharing-by-split show -node node1 -vserver vs1 -volume vol_clone1 -instance
Node Name: node1
Vserver Name: vs1
Volume Name: vol_clone1
Aggregate Name: aggr1

```
[[IDbeae1111dd7ccc79a163fc055194ba9e]]
= volume clone sharing-by-split undo show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the status of volume clone undo-sharing operations in-progress

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description
```

```
The `volume clone sharing-by-split undo show` command displays the
progress information of undo-sharing in the split volumes with shared
physical blocks. This command is only supported for flexible volumes. By
default, the command displays the following information:

* Vserver name
* Volume name
* Total number of blocks scanned for undo sharing
* Total number of blocks present
* Percentage of blocks processed



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This parameter selects information about the ongoing undo-sharing scan for
all volumes on this Vserver.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This parameter selects information about the ongoing undo-sharing scan on
this volume.

``[-blocks-scanned <integer>]`` - Scanned Blocks [.advaced-
priv]#(privilege: advanced)#::
This parameter selects information about the total number of blocks
scanned by undo-sharing in the given volume.

``[-blocks-total <integer>]`` - Total Blocks [.advaced-priv]#(privilege:
advanced)#::
This parameter selects information about the total number of blocks for
the undo-sharing to scan in the given volume.

``[-blocks-percentage-complete <integer>]`` - Blocks Percentage Complete
[.advaced-priv]#(privilege: advanced)#::
This parameter selects information about the percentage of block
```

```
processing completed by undo-sharing in the given volume.

== Examples

The following example displays information about all the ongoing undo-
sharing scan in the cluster:
```

cluster1::> volume clone sharing-by-split undo show
Blocks Blocks Blocks
Vserver Volume Scanned Total % Complete
--------- ---------- ---------- ---------- ----------
vs1 vol_clone1 0 1260 0

```
The following example displays information about volume ``_vol_clone1_``
residing on vserver ``_vs1_`` :
```

cluster1::> volume clone sharing-by-split undo show -vserver vs1 -volume vol_clone1 -instance
Vserver Name: vs1
Volume Name: vol_clone1
Blocks Scanned: 0
Blocks Total: 1260
Block Percentage complete: 0

```
[[ID3764f513df6b9a75b45f436c1ed2516e]]
= volume clone sharing-by-split undo start-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Undo the physical block sharing in split FlexClone volumes across the
cluster

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume clone sharing-by-split undo start-all` command starts a scan
to undo the shrared physical blocks in all the relevant volumes across the
cluster. The volumes will be available for the duration of the undo-
sharing operation. You can monitor the current progress of the scan using
the xref:{relative_path}volume-clone-sharing-by-split-undo-
show.html[volume clone sharing-by-split undo show] command. This command
is supported for flexible volumes that were split from their parent
volumes.

== Examples

The following example starts the scan to undo the physical block sharing
in all volumes across the cluster:
```

cluster1::> volume clone sharing-by-split undo start-all

```
== Related Links

* xref:{relative_path}volume-clone-sharing-by-split-undo-show.html[volume
clone sharing-by-split undo show]




[[ID78860ea6e659da377849c7edafccd56a]]
= volume clone sharing-by-split undo start
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Undo the physical block sharing in split FlexClone volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume clone sharing-by-split undo start` command starts a scan to
undo the shrared physical blocks in the given volume. The volume will be
available for the duration of the undo-sharing operation. After the scan
starts, you can stop it using the xref:{relative_path}volume-clone-
sharing-by-split-undo-stop.html[volume clone sharing-by-split undo stop]
command. You can monitor the current progress of the scan using the
xref:{relative_path}volume-clone-sharing-by-split-undo-show.html[volume
clone sharing-by-split undo show] command. This command is supported for
flexible volumes that were split from their parent volumes.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the vserver that the volume exists on.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the split volume with shared physical blocks, in
which the sharing will be undone.

== Examples

The following example starts the scan to undo the physical block sharing
in volume ``_vol_clone1_`` on vserver ``_vs1_`` :
```

cluster1::> volume clone sharing-by-split undo start -vserver vs1 -volume vol_clone1

```
== Related Links

* xref:{relative_path}volume-clone-sharing-by-split-undo-stop.html[volume
clone sharing-by-split undo stop]
* xref:{relative_path}volume-clone-sharing-by-split-undo-show.html[volume
clone sharing-by-split undo show]
```

```
[[IDf91d9d90b9c34dac34706bf10478a560]]
= volume clone sharing-by-split undo stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop an ongoing undo-sharing operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume clone sharing-by-split undo stop` command stops the process of
reverting the shared physical blocks from the split volume. If you restart
the undo-sharing operation, scan begins from the beginning because no
information about previously achieved progress is saved, but previously
unshared blocks are not processed again. This command is only supported
for flexible volumes.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the vserver that the volume exists on.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the volume whose unsharing of blocks will be
stopped.

== Examples

The following example stops an ongoing undo-sharing scan for volume
``_vol_clone1_`` on vserver ``_vs1_`` :
```

cluster1::> volume clone sharing-by-split undo stop -vserver vs1 -volume vol_clone1

```
[[IDbe1e3ff9b930f46f4820936207f62727]]
= volume clone split estimate
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Estimates the space required by the containing-aggregate to split the
FlexClone volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume clone split estimate` command displays an estimate of the free
disk space required in the aggregate to split the indicated clone volume
from its underlying parent volume. The value reported might differ from
the space actually required to perform the split, especially if the clone
volume is changing when the split is being performed. This command is only
supported for flexible volumes.

== Parameters

``[-vserver <vserver name>]`` - Vserver Name::
This specifies the estimates for free disk space required for splitting
FlexClone volumes residing on this Vserver. If the `-flexclone` option is
also specified, then the command displays the free disk space estimate
only for the specified FlexClone volume residing on the specified Vserver.

``[-flexclone <volume name>]`` - FlexClone Volume::
This specifies the free disk space estimate for splitting this FlexClone
volume.

``[-type {RW|DP}]`` - FlexClone Type::
This parameter specifies the type of FlexClone volume. A read-only
FlexClone volume is created if you specify the ``_type_`` as DP; otherwise
a read-write FlexClone volume is created.

``[-parent-vserver <vserver name>]`` - FlexClone Parent Vserver::
This specifies the free disk space estimates for splitting the FlexClone
volumes that are clones in the specified parent Vserver.

``[-b, -parent-volume <volume name>]`` - FlexClone Parent Volume::
This specifies the free disk space estimates for splitting the FlexClone
volumes cloned off this parent volume.

``[-parent-snapshot <snapshot name>]`` - FlexClone Parent Snapshot::
This specifies the free disk space estimates for splitting the FlexClone
volumes cloned off this parent snapshot.
```

```
``[-state {online|restricted|offline|force-online|force-offline|mixed}]``
- FlexClone Volume State::
This specifies the free disk space estimates for splitting the FlexClone
volumes with the specified state.

``[-junction-path <junction path>]`` - Junction Path::
This specifies the free disk space estimates for splitting the FlexClone
volumes mounted at this junction path.

``[-junction-active {true|false}]`` - Junction Active::
If this specified, the command displays the free disk space estimate for
splitting the FlexClone volumes with the specified junction path status.

``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
This specifies the free disk space estimates for splitting the FlexClone
volumes with the specified type of space guarantee.

``[-space-guarantee-enabled {true|false}]`` - Space Guarantee In Effect::
This specifies the free disk space estimates for splitting the FlexClone
volumes with the specified state of space guarantee.

``[-aggregate <aggregate name>]`` - FlexClone Aggregate::
This specifies the free disk space estimates for splitting the FlexClone
volumes residing on the specified aggregate.

``[-dsid <integer>]`` - FlexClone Data Set ID::
This specifies the free disk space estimates for splitting the FlexClone
volume with the specified DSID (data set ID).

``[-msid <integer>]`` - FlexClone Master Data Set ID::
This specifies the free disk space estimates for splitting the FlexClone
volumes with the specified MSID (master data set ID).

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - FlexClone Size::
This specifies the free disk space estimates for splitting FlexClone
volumes with the specified size.

``[-used {<integer>[KB|MB|GB|TB|PB]}]`` - Used Size::
This specifies the free disk space estimates for splitting the FlexClone
volumes with the specified amount of used disk space.

``[-split-estimate {<integer>[KB|MB|GB|TB|PB]}]`` - Split Estimate::
This specifies the free disk space estimates for splitting the FlexClone
volumes which match with the specified free disk space estimate for
splitting.
```

```
``[-blocks-scanned <integer>]`` - Blocks Scanned::
This specifies the free disk space estimates for splitting the FlexClone
volumes for which the specified number of blocks have been scanned.

``[-blocks-updated <integer>]`` - Blocks Updated::
This specifies the free disk space estimates for splitting the FlexClone
volumes for which the specified number of blocks have been updated.

``[-comment <text>]`` - Comment::
This specifies the free disk space estimates for splitting the FlexClone
volumes that have the specified value for the comment field.

``[-qos-policy-group <text>]`` - QoS Policy Group Name::
 This parameter optionally specifies which QoS policy group to apply to
the FlexClone volume. The policy group defines measurable service level
objectives (SLOs) that apply to the storage objects with which the policy
group is associated. If you do not assign a policy group to the FlexClone
volume, the system does not monitor and control the traffic to the volume.

``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name::
 This optionally specifies which QoS adaptive policy group to apply to the
volume. This policy group defines measurable service level objectives
(SLOs) and Service Level agreements (SLAs) that adjust based on the volume
allocated space or used space. This parameter is not supported on
FlexGroup volumes.

``[-caching-policy <text>]`` - Caching Policy Name::
This specifies the free disk space estimates for splitting the FlexClone
volumes that are clones with the specified caching policy.

``[-parent-vol-type <volAccessType>]`` - Parent volume type [.advaced-
priv]#(privilege: advanced)#::
This specifies the free disk space estimates for splitting the FlexClone
volumes that are clones of the parent volumes with the specified type.

``[-flexclone-used-percent <percent>]`` - FlexClone Used Percentage::
This specifies the free disk space estimates for splitting the FlexClone
volumes that are clones with the specified percentage of used space.

``[-vserver-dr-protection {protected|unprotected}]`` - Vserver DR
Protection::
This specifies the free disk space estimates for splitting the FlexClone
volumes that are clones with the specified Vserver SnapMirror protection.

``[-block-percentage-complete <integer>]`` - Percentage Complete::
This specifies the free disk space estimates for splitting the FlexClone
```

```
volumes for which the specified percentage of Block processing has been
completed.

``[-uid <integer>]`` - Volume-Level UID::
This specifies the free disk space estimates for splitting the FlexClone
volumes that are clones with the specified volume-level UID.

``[-gid <integer>]`` - Volume-Level GID::
This specifies the free disk space estimates for splitting the FlexClone
volumes that are clones with the specified volume-level GID.

``[-flexgroup-uuid <UUID>]`` - UUID of the FlexGroup::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volume constituents that have the
specified FlexGroup volume UUID. This parameter applies to FlexGroup
volumes and FlexGroup volume constituents only.

``[-flexgroup-msid <integer>]`` - FlexGroup Master Data Set ID::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volume constituents that have the
specified FlexGroup volume master data-set ID. This parameter applies to
FlexGroup volumes and FlexGroup volume constituents only.

``[-flexgroup-index <integer>]`` - FlexGroup Index::
If this parameter is specified, the command displays information only
about the FlexGroup volume constituents that have the specified FlexGroup
volume index. This parameter applies to FlexGroup volume constituents
only.

``[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum size of
a FlexGroup Constituent::
If this parameter is specified, the command displays information only
about the FlexGroup volume or FlexGroup volumes that have the specified
maximum constituent size. This parameter applies to FlexGroup volumes
only.

``[-constituent-role <Constituent Roles>]`` - Constituent Volume Role::
If this parameter is specified, the command displays information only
about the constituent volume or volumes that are of the specified role.
This parameter applies to FlexGroup volumes only.

``[-is-flexgroup-constituent-active {true|false}]`` - Is Active FlexGroup
Constituent::
If this parameter is specified, the command displays information only
about the volume or volumes that are either active constituent of the
FlexGroup or not, depending on the value provided.
```

```
``[-is-constituent {true|false}]`` - Is Constituent Volume::
If this parameter is specified, the command displays information only
about volumes that either are or are not constituents of a FlexGroup,
depending on the value provided.

``[-is-flexgroup {true|false}]`` - Is Volume a FlexGroup::
If this parameter is specified, the command displays information only
about the volume or volumes that are either FlexGroup volumes or not,
depending on the value provided.

``[-volume-style-extended {flexvol|flexgroup|flexgroup-constituent}]`` -
Extended Volume Style::
If this parameter is specified, the command displays information only
about the volumes that are configured with the specified extended style.
Possible values are ``_flexvol_`` for FlexVol volumes, ``_flexgroup_`` for
FlexGroup volumes and ``_flexgroup-constituent_`` for FlexGroup
constituents.

``[-snaplock-type {non-snaplock|compliance|enterprise}]`` - SnapLock
Type::
If this parameter is specified, the command displays information only
about the volumes that are configured with the specified SnapLock type.
Possible values are ``_compliance_`` for SnapLock Compliance volumes,
``_enterprise_`` for SnapLock Enterprise volumes and ``_non-snaplock_``
for non-SnapLock volumes.

== Examples

The following example displays the FlexClone split free disk space
estimates for the FlexClone volumes residing on Vserver ``_vs0_`` .
```

cluster1::> volume clone split estimate -vserver vs0
(volume clone split estimate)
Split
Vserver FlexClone Estimate
--------- ------------ ----------
vs0 fc_vol_1 851.5MB
fc_vol_3 0.00B
flex_clone1 350.3MB
fv_2 47.00MB
tv9 0.00B
5 entries were displayed.

```
[[IDe2cae2bf68ec28ef904a4394c31047df]]
= volume clone split show
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the status of FlexClone split operations in-progress

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume clone split show` command displays the progress information of
all the active FlexClone volume splitting jobs. This command is only
supported for flexible volumes. By default, this command displays the
following information:

* Vserver name
* FlexClone volume name
* Percentage of blocks processed
* Total number of blocks scanned for clone splitting
* Total number of blocks updated for clone splitting

If the -instance option is also specified, detailed information about all
splitting jobs is displayed.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields to be displayed, for all the ongoing FlexClone
splitting jobs.

| ``[-instance  ]`` }::
This specifies the command to display detailed information about the
ongoing FlexClone volume splitting jobs.

``[-vserver <vserver name>]`` - Vserver Name::
Selects information about the ongoing FlexClone volume splitting jobs for
all FlexClone volumes on this Vserver.

``[-flexclone <volume name>]`` - FlexClone Volume::
Selects information about ongoing FlexClone volume splitting jobs for this
FlexClone volume.

``[-block-percentage-complete <integer>]`` - Percentage Complete::
Selects information about all the ongoing FlexClone splitting jobs that
have the specified percentage of Block processing completed.
```

```
``[-blocks-scanned <integer>]`` - Blocks Scanned::
Selects information about all the ongoing FlexClone splitting jobs that
have the specified number of blocks scanned.

``[-blocks-updated <integer>]`` - Blocks Updated::
Selects information about all the ongoing FlexClone splitting jobs that
have the specified number of blocks updated.

== Examples

The following example displays information about all the ongoing FlexClone
splitting jobs in the cluster.
```

cluster1::> volume clone split show
(volume clone split show)
Blocks

Vserver FlexClone Scanned Updated % Complete

--- --- --- --- ---
vs1 fc_vol_1 229 2 0

```
The following example displays information about FlexClone volume
``_fc_vol_2_`` residing on Vserver ``_vs0_`` .
```

cluster1::> volume clone split show -vserver vs0 -flexclone fc_vol_2 -instance
(volume clone split show)
Vserver Name: vs0
FlexClone Volume: fc_vol_2
Percentage Complete: 0
Blocks Scanned: 229
Blocks Updated: 2

```
[[IDffc8d8988412f4e006cbc9df26abb044]]
= volume clone split start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Split a FlexClone from the parent volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `volume clone split start` command starts a job to separate the
FlexClone volume from the underlying parent volume. Both, the parent and
the FlexClone volumes will be available for the duration of the split
operation. After the job starts, you can stop it using the
xref:{relative_path}volume-clone-split-stop.html[volume clone split stop]
command. You can also stop the job using the xref:{relative_path}job-
stop.html[job stop] command. You can monitor the current progress of the
job using the xref:{relative_path}volume-clone-split-show.html[volume
clone split show] and xref:{relative_path}job-show.html[job show]
commands. This command is only supported for flexible volumes. This
command is not supported on volumes that are being protected as part of a
Vserver level SnapMirror.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver that the FlexClone volume exists on.

``-flexclone <volume name>`` - FlexClone Volume::
This specifies the FlexClone volume that will be split from its parent
volume.

``[-foreground <true>]`` - Foreground Process::
This specifies whether the clone splitting job will run as a foreground
job. The default value of this option is `true` .

``[-match-parent-storage-tier {true|false}]`` - Match Split Data to Parent
Storage Tier::
This specifies whether the FlexClone volume splits the data blocks by
matching its parent storage tier. This option is applicable only if the
tiering policy and the tiering minimum cooling days of the parent volume
and the FlexClone volume are the same. The default value of this option is
`false` .

== Examples

The following example starts splitting FlexClone volume ``_fc_vol_1_`` on
Vserver ``_vs1_`` as a foreground job.
```

cluster1::> volume clone split start -vserver vs1 -flexclone fc_vol_1 -foreground true

```
== Related Links

* xref:{relative_path}volume-clone-split-stop.html[volume clone split
```

```
stop]
* xref:{relative_path}job-stop.html[job stop]
* xref:{relative_path}volume-clone-split-show.html[volume clone split
show]
* xref:{relative_path}job-show.html[job show]




[[ID2b4ae07ae8b63fb018ce79dcba1a15cf]]
= volume clone split stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop an ongoing FlexClone split job

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume clone split stop` command stops the process of separating the
FlexClone volume from its underlying parent volume, but does not lose any
of the progress achieved while the split process was active. That is, all
the clone volume blocks already separated from the parent volume remain
separated. If you restart the split operation, splitting process begins
from the beginning because no information about previously achieved
progress is saved, but previously split blocks are not re-split. This
command is only supported for flexible volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver that the FlexClone volume exists on.

``-flexclone <volume name>`` - FlexClone Volume::
This specifies the FlexClone volume whose separation from the parent
volume will be stopped.

== Examples

The following example stops an ongoing clone splitting job for FlexClone
volume ``_fc_vol_1_`` on Vserver ``_vs1_`` .
```

cluster1::> volume clone split stop -vserver vs1 -flexclone fc_vol_1

```
:leveloffset: -1


= volume conversion commands

:leveloffset: +1


[[ID43edb98ef1f7a19b9cc6fdfa8bc55c26]]
= volume conversion start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Convert a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume conversion start` command converts a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver in which the volume is located.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume to convert.

``[-check-only <true>]`` - Validate the Conversion Only [.advaced-
priv]#(privilege: advanced)#::
If true, this specifies whether the command will only validate that the
conversion can be done, and that the volume will not be converted. The
default value is false.

``[-foreground <true>]`` - Foreground Process [.advaced-priv]#(privilege:
advanced)#::
This specifies whether the operation runs in the foreground. The default
```

```
setting is ``_true_`` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes.

== Examples

The following example validates and converts a flexible volume named
"flexvol" in Vserver "vs1" to a FlexGroup:
```

cluster::*> volume conversion start -vserver vs1 -volume flexvol

```
The following example validates whether flexible volume named "flexvol" in
Vserver "vs1" can be converted to a FlexGroup:
```

cluster::*> volume conversion start -vserver vs1 -volume flexvol -check-only

```
[[IDda7d7c1b551c559efad18b92bf54d5d6]]
= volume conversion validation show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the result of volume conversion pre checks.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume conversion validation show` command displays the results of
volume conversion validation for volumes.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
```

```
advanced)#::
If this parameter is specified, the command displays information only
about the conversion operations that are running on volumes on the
specified Vserver.

``[-volume <volume name>]`` - Volume [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the conversion operations on the specified volume.

``[-precheck <Volume conversions pre check>]`` - Volume Feature [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the specified conversion pre checks.

``[-status <Feature status for volume conversion>]`` - Feature Status
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the conversion pre checks with the specified status. The supported
statuses are:
+
  * supported -The volume can be converted.
  * warn - The volume can be converted, though there may be an issue
highlighted by a warning in the `error` value.
  * block - The volume cannot be converted until the specified `action` is
taken
  * unsupported - The volume cannot be converted.
  * error - There was an error performing the pre check.



``[-error <text>]`` - Feature Error [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the conversion pre checks with the specified error message.

``[-action <text>]`` - Feature Action [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the onversion pre checks with the specified resolution action.

== Examples

The following example displays all conversion pre checks
```

cluster::*> volume conversion validation show

```
The following example displays all conversion pre checks that block
conversion until an action is taken, for volume 'vol1' in Vserver 'vs0'
```

cluster::*> volume conversion validation show -vserver vs0 -volume vol1 -status block

```
:leveloffset: -1


= volume efficiency commands

:leveloffset: +1


[[ID d1e10e0575451f9709d9fcd3649f8923]]
= volume efficiency check
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Scrub efficiency metadata of a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command verifies and updates the fingerprint database for the
specified volume. This command is not supported on FlexGroup volumes or
Infinite Volumes that are managed by storage services.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the volume on which the verify operation needs to be started.

| ``-path </vol/volume>`` - Volume Path [.advaced-priv]#(privilege:
advanced)# }::
Specifies the volume path on which the verify operation needs to be
```

```
started.

``[-d, -delete-checkpoint {true|false}]`` - Delete Checkpoint [.advaced-
priv]#(privilege: advanced)#::
Deletes existing checkpoint.

== Examples

The following example runs `volume efficiency check` with delete
checkpoint option turned on.
```

cluster1::*> volume efficiency check -vserver vs1 -volume vol1 -delete-checkpoint true

```
[[ID78e8bfe55da686aabfdde668ee1646f1]]
= volume efficiency modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the efficiency configuration of a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name::
This specifies the volume on which efficiency options need to be modified.

| ``-path </vol/volume>`` - Volume Path }::
This specifies the volume path on which efficiency options need to be
modified.

{ ``[-schedule <text>]`` - Schedule::
This option is used to set and modify the schedule.
+
schedule is [day_list][@hour_list] or [hour_list][@day_list] or - or auto
or manual
+
The day_list specifies the days of the week that an efficiency operation
should run. It is a list of the first three letters of the day (sun, mon,
```

```
tue, wed, thu, fri, sat), separated by a comma. Day ranges such as mon-fri
can also be used. The default day_list is sun-sat. The names are not case
sensitive.
+
The hour_list specifies the hours of each scheduled day that an efficiency
operation should run. The hour_list is from 0 to 23, separated by a comma.
Hour ranges such as 8-17 are allowed. Step values can be used in
conjunction with ranges (For example, 0-23/2 means every two hours in a
day). The default hour_list is 0, i.e. at midnight of each scheduled day.
+
When efficiency is enabled on a volume for the first time, an initial
schedule is assigned to the volume. This initial schedule is sun-sat@0,
which means run once every day at midnight.
+
If "-" is specified, no schedule is set on the volume. The auto schedule
string triggers an efficiency operation depending on the amount of new
data written to the volume. The manual schedule string prevents SIS from
automatically triggering any operations and disables change-logging. This
schedule string can only be used on SnapVault destination volumes. The use
of this schedule is mainly desirable when inline compression is enabled on
a SnapVault destination volume and background processing is not necessary.
+
Note that schedule and policy are mutually exclusive options.

| ``[-policy <text>]`` - Efficiency Policy Name }::
This option is used to set an efficiency policy. The policy cannot be
changed to the predefined ``_inline-only_`` and ``_auto_`` policy when
there is an active background operation on the volume. The policies
``_inline-only_`` and ``_none_`` are not supported on Capacity optimized
Flash with QAT supported platforms.
+
Note that schedule and policy are mutually exclusive options.

{ ``[-compression {true|false}]`` - Compression::
This option is used to enable and disable compression. The default value
is determined based on the platform.

``[-inline-compression {true|false}]`` - Inline Compression::
This option is used to enable and disable inline compression. Inline
compression can be enabled only if `compression` is enabled. The default
value is determined based on the platform.
+
You can use the ``_inline-only_`` predefined efficiency policy to run
inline compression without the need of any background efficiency
operations.
```

```
``[-compression-type {none|secondary|adaptive}]`` - Compression Type
[.advaced-priv]#(privilege: advanced)#::
This option is used to specify the size of compression group on the
volume. The default value is determined based on the platform.

| ``[-storage-efficiency-mode {default|efficient}]`` - Storage Efficiency
Mode }::
This option is used to modify the storage efficiency mode. The default
mode sets 8k adaptive compression on the volume. The efficient mode sets
auto adaptive compression and attempts to enable inline deduplication,
cross volume inline deduplication and cross volume background
deduplication on the volume. There is no difference between default and
efficient modes on QAT supported platforms and auto adaptive compression
is set irrespective of the modes.
+
The available efficiency modes are:
+
  * default
  * efficient



``[-inline-dedupe {true|false}]`` - Inline Dedupe::
This option is used to enable and disable inline deduplication. The
default value is determined based on the platform.
+
You can use the ``_inline-only_`` predefined efficiency policy to run
inline deduplication without the need of any background efficiency
operations.

``[-data-compaction {true|false}]`` - Data Compaction::
This option is used to enable and disable data compaction. The default
value is determined based on the platform.

``[-cross-volume-inline-dedupe {true|false}]`` - Cross Volume Inline
Deduplication::
This option is used to enable and disable cross volume inline
deduplication. The default value is determined based on the platform.

``[-cross-volume-background-dedupe {true|false}]`` - Cross Volume
Background Deduplication::
This option is used to enable and disable cross volume background
deduplication. The default value is determined based on the platform.

== Examples

The following examples modify efficiency options on a volume.
```

```
cluster1::> volume efficiency modify -vserver vs1 -volume vol1 -schedule sun-sat@12
```

```

```

cluster1::> volume efficiency modify -vserver vs1 -volume vol1 -policy policy1

```

```

cluster1::> volume efficiency modify -vserver vs1 -volume vol1 -compression true -inline-compression true -inline-dedupe true -data-compaction true -cross-volume-inline-dedupe true -cross-volume-background-dedupe true

```
[[IDf02fae37b4bd46acb17fba890b84d4cf]]
= volume efficiency off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disables efficiency on a volume

*Availability:* This command is available to _cluster_ and _Vserver_ administrators at the _admin_ privilege level.

== Description

The `volume efficiency off` command disables efficiency on a volume. Disabling efficiency is not supported on Capacity optimized Flash with QAT supported platforms.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name::
Specifies the name of the volume on which efficiency needs to be disabled.

| ``-path </vol/volume>`` - Volume Path }::
Specifies the volume path on which efficiency needs to be disabled.

== Examples

The following examples disable efficiency on a volume:
```

cluster1::> volume efficiency off -vserver vs1 -volume vol1

```
```

cluster1::> volume efficiency off -vserver vs1 -path /vol/vol1

```
[[ID3c847770681d4bbbc298de856eb1bd73]]
= volume efficiency on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable efficiency on a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume efficiency on` command enables efficiency on a volume. The
specified volume must be online. Efficiency operations will be started
periodically according to a per volume schedule or policy. The
xref:{relative_path}volume-efficiency-modify.html[volume efficiency
modify] command can be used to modify schedule and the
xref:{relative_path}volume-efficiency-policy-modify.html[volume efficiency
policy modify] command can be used to modify policy. You can also manually
start an efficiency operation with the xref:{relative_path}volume-
efficiency-start.html[volume efficiency start] command.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which efficiency needs to be
enabled.

| ``-path </vol/volume>`` - Volume Path }::
This specifies the volume path on which efficiency needs to be enabled.

== Examples

The following examples enable efficiency on a volume.
```

cluster1::> volume efficiency on -vserver vs1 -volume vol1

```

```

cluster1::> volume efficiency on -vserver vs1 -path /vol/vol1

```
== Related Links

* xref:{relative_path}volume-efficiency-modify.html[volume efficiency
modify]
* xref:{relative_path}volume-efficiency-policy-modify.html[volume
efficiency policy modify]
* xref:{relative_path}volume-efficiency-start.html[volume efficiency
start]




[[IDec061360ed730857bec354535d53e83d]]
= volume efficiency prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Identify any incompatible volumes or Snapshot copies before downgrade

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency prepare-to-downgrade` command updates efficiency
configurations and metadata to be compatible with prior releases as
specified. This command also disables the use of incompatible efficiency
features. This command is not supported on FlexGroup volumes.

== Parameters

``[-disable-feature-set <downgrade version>]`` - Data ONTAP Version
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the ONTAP version that introduced new volume
efficiency feature set.

== Examples

The following example disables the features introduced in ONTAP 8.3.1
```

cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set 8.3.1

```
The following example disables the features introduced in ONTAP 9.6.0.
```

cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set 9.6.0

```
The following example ignores offline volumes while disabling the features
introduced in ONTAP 9.6.0 .
```

cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set 9.6.0 -skip-offline-volumes true

```
The following example ignores offline volumes while disabling the features
introduced in ONTAP 8.3.1 .
```

cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set 8.3.1 -skip-offline-volumes true

```
[[IDe78f63da5332c2b051928cd14f1786a8]]
= volume efficiency promote
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a volume to the preferred set of volumes for efficiency processing

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

Use the `volume efficiency promote` command to promote a volume from
deprioritized state back to auto state.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume on which auto scheduling needs to be
restarted.

| ``-path </vol/volume>`` - Volume Path [.advaced-priv]#(privilege:
advanced)# }::
This specifies the volume path on which auto scheduling needs to be
restarted.

== Examples

The following example promotes a volume from deprioritized state back to
auto state.
```

cluster1::*> volume efficiency promote -vserver vs1 -volume vol1

```
[[ID19c17b7f9f600a2c9560327b6f9b8b52]]
= volume efficiency revert-to
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reverts volume efficiency metadata

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency revert-to` command reverts the format of volume
efficiency metadata for the volume to the given version of ONTAP. This
command is not supported on FlexGroup volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume for which volume efficiency metadata
needs to be reverted.

| ``-path </vol/volume>`` - Volume Path [.advaced-priv]#(privilege:
advanced)# }::
This specifies the volume path for which volume efficiency metadata needs
to be reverted.

``[-version <revert version>]`` - Revert to Version [.advaced-
priv]#(privilege: advanced)#::
Specifies the version of ONTAP to which the volume efficiency metadata
needs to be formatted.

``[-d, -delete {true|false}]`` - Delete Existing Metafile on Revert
[.advaced-priv]#(privilege: advanced)#::
If set to `true` , this parameter specifies that the volume efficiency
metadata be deleted instead of reverting its format. By default this
parameter is set to `false` .

``[-c, -clean-up {true|false}]`` - Delete Previously Downgraded Metafiles
[.advaced-priv]#(privilege: advanced)#::
If set to `true` , this parameter specifies that the volume efficiency
metadata already reverted using `volume efficiency revert-to ` be deleted.
By default this parameter is set to `false` .
```

```
``[-revert-adaptive-compression {true|false}]`` - Downgrade to minor
version [.advaced-priv]#(privilege: advanced)#::
If set to `true` , this parameter specifies that the volume efficiency
metadata needs to be reverted to a minor version of ONTAP. By default this
parameter is set to `false` .

``[-check-snapshot {true|false}]`` - Revert ignore snapshots [.advaced-
priv]#(privilege: advanced)#::
If set to ``_false_`` , this parameter specifies that the volume
efficiency revert will not check for snapshots created by previous
releases of ONTAP. By default this parameter is set to ``_true_`` .

== Examples

The following examples reverts volume efficiency metadata on a volume
named vol1 located in vserver vs1 to version 9.8.
```

cluster1::*> volume efficiency revert-to -vserver vs1 -volume vol1 -version 9.8

```

```

cluster1::*> volume efficiency revert-to -vserver vs1 -path /vol/vol1 -version 9.8

```
[[IDb4082219a147d59217bb0db6a6ff4ffa]]
= volume efficiency show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of volumes with efficiency

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume efficiency show` command displays the information about
storage efficiency of volumes. The command output depends on the parameter
or parameters specified. If no parameters are specified, the command
displays the following information for all volumes with efficiency:

* Vserver: Vserver the volume belongs to.
* Volume: Name of the volume.
```

```
* State: Current state of efficiency on the volume (Enabled, Disabled, or
Mixed).
* Status: Status of the efficiency on the volume. Following are the
possible values:

  * ``_Active_`` : An efficiency operation is currently running.
  * ``_Idle_`` : There are no efficiency operations running.
  * ``_Initializing_`` : An efficiency operation is being initialized.
  * ``_Undoing_`` : Efficiency is being undone on the volume.
  * ``_Pending_`` : An efficiency operation is queued.
  * ``_Downgrading_`` : An efficiency operation necessary to downgrade the
efficiency metafiles to a previous ONTAP release is active.
  * ``_Disabled_`` : Efficiency is disabled on the volume.


* Progress: The progress of the current efficiency operation with
information as to which stage of the efficiency process is currently in
progress and how much data is processed for that stage. For example: "25
MB Scanned", "20 MB Searched", "500 KB (2%) Compressed", "40 MB (20%)
Done", "30 MB Verified".



To display detailed information, run the command with the ` -l ` or `-
instance` parameter. The detailed view provides all information in the
previous list and the following additional information:

* Path: Volume Path.
* Compression: Current state of compression on the volume (Enabled or
Disabled).
* Inline Compression: Current state of inline compression on the volume
(Enabled or Disabled).
* Type: Type of volume (Regular or SnapVault).
* Schedule: The schedule of efficiency operation for the volume.
* Policy: Efficiency policy for the volume.
* Minimum Blocks Shared: The minimum number of adjacent blocks in a file
that can be shared.
* Blocks Skipped Sharing: Blocks skipped sharing because of the minimum
block share value.
* Last Operation State: Status of the last operation (Success or Failure).
* Last Successful Operation Begin: The time and date at which the last
successful operation began.
* Last Successful Operation End: The time and date at which the last
successful operation ended.
* Last Operation Begin: The time and date at which the last operation
began.
```

```
* Last Operation End: The time and date at which the last operation ended.
* Last Operation Size: The size of the last operation.
* Last Operation Error: The error encountered by the last operation.
* Change Log Usage: The percentage of the change log that is used.
* Logical Data: The total logical data in the volume, and how much is
reached compared to the deduplication logical data limit.
* Queued Job: The job that is queued. Following are the possible values:

  * ``_-_`` : There are no queued jobs.
  * ``_scan_`` : A job to process existing data is queued.
  * ``_start_`` : A job to process newly added data is queued.
  * ``_check_`` : A job to eliminate stale data from the fingerprint
database is queued.
  * ``_downgrading_`` : An efficiency operation necessary to downgrade the
efficiency metafiles to a previous ONTAP release is queued.


* Stale Fingerprints: The percentage of stale entries in the fingerprint
database. If this is greater than 20 percent a subsequent
xref:{relative_path}volume-efficiency-start.html[volume efficiency start]
operation triggers the verify operation, which might take a long time to
complete.
* Inline Dedupe: Current state of inline deduplication on the volume
(Enabled or Disabled).
* Cross Volume Inline Deduplication: Current state of cross volume inline
deduplication on the volume (Enabled or Disabled).
* Cross Volume Background Deduplication: Current state of cross volume
background deduplication on the volume (Enabled or Disabled).
* Extended Compressed Data: Is there extended compressed data present on
the volume.
* Inline Adaptive Data Compaction: Whether Inline Adaptive Data Compaction
is enabled or disabled on the volume. When enabled, ONTAP combines data
fragments to reduce on-disk block consumption.



You can specify additional parameters to display information that matches
only those parameters. For example, to display information only about
volumes with efficiency in Vserver vs1, run the command with the `-vserver
vs1` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed. The fields Vserver
and volume name are the default fields.
```

```
| ``[-l  ]``::
This option displays detailed information about the volumes with
efficiency.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name::
Displays information only for those volumes that match the specified
Vserver.

{ ``[-volume <volume name>]`` - Volume Name::
Displays information only for those volumes that match the specified
volume.

| ``[-path </vol/volume>]`` - Volume Path }::
Displays information only for those volumes that match the specified
volume path.

``[-state {Disabled|Enabled|Mixed}]`` - State::
Displays information only for those volumes that match the specified
state.

``[-op-status <Efficiency status>]`` - Status::
Displays information only for those volumes that match the specified
operation status.

``[-progress <text>]`` - Progress::
Displays information only for those volumes that match the specified
progress.

``[-type {Regular|SnapVault}]`` - Type::
Displays information only for those volumes that match the specified type
of volume.

``[-schedule <text>]`` - Schedule::
Displays information only for those volumes that match the specified
schedule.

``[-policy <text>]`` - Efficiency Policy Name::
Displays information only for those volumes that match the specified
policy.

``[-blks-skipped-sharing <integer>]`` - Blocks Skipped Sharing::
Displays information only for those volumes that match the specified
blocks skipped sharing.
```

```
``[-last-op-state <text>]`` - Last Operation State::
Displays information only for those volumes that match the specified last
operation state.

``[-last-success-op-begin <Date>]`` - Last Success Operation Begin::
Displays information only for those volumes that match the specified last
successful operation begin time.

``[-last-success-op-end <Date>]`` - Last Success Operation End::
Displays information only for those volumes that match the specified last
successful operation end time.

``[-last-op-begin <Date>]`` - Last Operation Begin::
Displays information only for those volumes that match the specified last
operation begin time.

``[-last-op-end <Date>]`` - Last Operation End::
Displays information only for those volumes that match the specified last
operation end time.

``[-last-op-size {<integer>[KB|MB|GB|TB|PB]}]`` - Last Operation Size::
Displays information only for those volumes that match the specified last
operation size.

``[-last-op-error <text>]`` - Last Operation Error::
Displays information only for those volumes that match the specified last
operation error.

``[-changelog-usage <percent_no_limit>]`` - Changelog Usage::
Displays information only for those volumes that match the specified
change log usage.

``[-logical-data-size {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Data Size::
Displays information only for those volumes that match the specified
logical data size.

``[-logical-data-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Logical Data
Limit::
Displays information only for those volumes that match the specified
logical data limit.

``[-logical-data-percent <percent_no_limit>]`` - Logical Data Percent::
Displays information only for those volumes that match the specified
logical data percentage.

``[-queued-job <text>]`` - Queued Job::
```

```
Displays information only for those volumes that match the specified
number of queued jobs.

``[-stale-fingerprint-percentage <integer>]`` - Stale Fingerprint
Percentage::
Displays information only for those volumes that match the specified stale
fingerprint percentage.

``[-compression {true|false}]`` - Compression::
Displays information only for those volumes that match the specified
compression setting.

``[-inline-compression {true|false}]`` - Inline Compression::
Displays information only for those volumes that match the specified
inline compression setting.

``[-compression-type {none|secondary|adaptive}]`` - Compression Type
[.advaced-priv]#(privilege: advanced)#::
Displays information about the type of compression on the volume[adaptive
or secondary].

``[-storage-efficiency-mode {default|efficient}]`` - Storage Efficiency
Mode::
Displays information only for those volumes that match the specified
storage efficiency mode. The default mode sets 8k adaptive compression on
the volume. The efficient mode sets auto adaptive compression on the
volume.
+
The available efficiency modes are:
+
  * default
  * efficient



``[-is-constituent {true|false}]`` - Constituent Volume::
Displays information only for those volumes that either are or are not
constituents of a FlexGroup, depending on the value provided.

``[-inline-dedupe {true|false}]`` - Inline Dedupe::
Displays information only for those volumes that match the specified
inline deduplication setting.

``[-data-compaction {true|false}]`` - Data Compaction::
Displays information only for those volumes that match the specified data
compaction setting.
```

```
``[-cross-volume-inline-dedupe {true|false}]`` - Cross Volume Inline
Deduplication::
Displays information only for those volumes that match the specified cross
volume inline deduplication setting.

``[-cross-volume-background-dedupe {true|false}]`` - Cross Volume
Background Deduplication::
Displays information only for those volumes that match the specified cross
volume background deduplication setting.

``[-extended-compressed-data {true|false}]`` - Extended Compressed Data::
Displays information only for those volumes that match the specified
extended compressed data value. Extended compressed data is enabled on a
volume when both adaptive compression configured with application IO size
8K and data compaction are enabled. Once enabled, extended compressed data
can only be disabled by using the `volume efficiency revert-to` command.

``[-extended-auto-adaptive-compression {true|false}]`` - Volume Has
Extended Auto Adaptive Compression [.advaced-priv]#(privilege:
advanced)#::
Displays information only for those volumes that match the specified
extended auto adaptive compression setting.

``[-active-data-extended-auto-adaptive-compression {true|false}]`` -
Volume Has Active Data Extended Auto Adaptive Compression [.advaced-
priv]#(privilege: advanced)#::
Displays information only for those volumes that match the specified
active data extended auto adaptive compression setting.

== Examples

The following example displays information about all volumes with
efficiency on the Vserver named vs1:
```

```
cluster1::> volume efficiency show -vserver vs1
Vserver Volume State Status Progress
----------- ------------------ -------- ------------ ------------------
vs1 vol1 Enabled Idle Idle for 22:37:53
vs1 vol2 Enabled Idle Idle for 22:37:53
vs1 vol3 Enabled Idle Idle for 22:37:49
vs1 vol4 Enabled Idle Idle for 22:37:53
vs1 vol5 Enabled Idle Idle for 22:37:53
vs1 volham Enabled Idle Idle for 22:37:53
vs1 volham1 Enabled Idle Idle for 22:37:53
7 entries were displayed.
```

```
The following example displays detailed information about a volume named
vol1 on a Vserver named vs1:
```

cluster1::> volume efficiency show -vserver vs1 -volume vol1
Vserver Name: vs1
Volume Name: vol1
Volume Path: /vol/vol1
State: Enabled
Status: Idle
Progress: Idle for 02:14:28
Type: Regular
Schedule: sun-sat@0
Efficiency Policy Name: -
Blocks Skipped Sharing: 0
Last Operation State: Success
Last Success Operation Begin: Wed Jul 14 03:08:44 2021
Last Success Operation End: Wed Jul 14 03:08:46 2021
Last Operation Begin: Wed Jul 14 03:08:44 2021
Last Operation End: Wed Jul 14 03:08:46 2021
Last Operation Size: 2.52MB
Last Operation Error: -
Changelog Usage: 0%
Logical Data Size: 20.48MB
Logical Data Limit: 640TB
Logical Data Percent: 0%
Queued Job: -
Stale Fingerprint Percentage: 0
Compression: true
Inline Compression: true
Storage Efficiency Mode: -
Constituent Volume: false
Inline Dedupe: false
Data Compaction: false
Cross Volume Inline Deduplication: false
ross Volume Background Deduplication: false
Extended Compressed Data: false

```
== Related Links

* xref:{relative_path}volume-efficiency-start.html[volume efficiency
start]




[[ID8aa56aaa3b702fce971010013c9f124e]]
= volume efficiency start
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Starts efficiency operation on a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `volume efficiency start` command to start an efficiency
operation. The volume must be online and have efficiency enabled. If there
is an efficiency operation already active on the volume, this command
fails.

When the `volume efficiency start` command is issued, a checkpoint is
created at the end of each stage or sub-stage, or on an hourly basis in
the gathering phase. If at any point the `volume efficiency start`
operation is stopped, the system can restart the efficiency operation from
the execution state saved in the checkpoint. The `delete-checkpoint`
parameter can be used to delete the existing checkpoint and restart a
fresh efficiency operation. The checkpoint corresponding to gathering has
a validity period of 24 hours. If the user knows that significant changes
have not been made on the volume, then such a gatherer checkpoint whose
validity has expired can be used with the help of the `use-checkpoint`
parameter. There is no time restriction for checkpoints of other stages.

When the volume is configured to use the ``_inline-only_`` efficiency
policy, the system will stop monitoring changes to the data for the
purpose of running background efficiency operations. The background
deduplication operations will be disabled. The user can still execute
compression specific efficiency operation with `-scan-old-data` and `-
compression` parameters to compress the existing data on the volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name::
Specifies the name of the volume.

| ``-path </vol/volume>`` - Volume Path }::
Specifies the complete path of the volume.

``[-s, -scan-old-data <true>]`` - Scan Old Data::
```

```
This option scans the file system and processes all existing data. It
prompts for user confirmation before proceeding. Use the force option to
suppress this confirmation.

{ ``[-p, -use-checkpoint <true>]`` - Use Checkpoint (if scanning old
data)::
Use the checkpoint when scanning existing data. Valid only if `scan-old-
data` parameter is true.

| ``[-d, -delete-checkpoint <true>]`` - Delete Checkpoint }::
Deletes the existing checkpoint and restarts a new `volume efficiency
start` operation.

``[-qos-policy <sis_qos>]`` - QoS Policy::
Specifies the qos-policy, which indicates how the efficiency operations
are throttled. This option can be configured to be `background` or `best-
effort` . Default value is `best-effort` . If `background` is specified,
the efficiency operations are run with minimum or no impact on the data
serving client operations. If `best-effort` is specified, the efficiency
operations might have some impact on the data serving client operations.

``[-C, -compression <true>]`` - Start Compression (if scanning old data)
[.advaced-priv]#(privilege: advanced)#::
Compresses existing data. Deduplication is not run unless the dedupe
option is also specified. Valid only if `scan-old-data` parameter is true.

``[-D, -dedupe <true>]`` - Start Deduplication (if scanning old data)
[.advaced-priv]#(privilege: advanced)#::
Deduplicates existing data on disk. Similarly, compression is not run
unless the compression option is also specified. Valid only if `scan-old-
data` parameter is true.

``[-P, -compaction <true>]`` - Start Compaction (if scanning old data)
[.advaced-priv]#(privilege: advanced)#::
Compacts existing data on disk. Valid only if `scan-old-data` parameter is
true.

``[-m, -build-metadata <true>]`` - Build metadata without sharing(if
scanning old data)::
Builds deduplication metadata by scanning the entire file system. You will
not achieve any space savings with this option. Once the metadata is
built, existing data can be shared with newly written data on subsequent
deduplication runs.

``[-o, -scan-all <true>]`` - Scan all the data without shared block
optimization(if scanning old data)::
```

```
Scans the entire file system and processes the shared blocks also. You may
be able to achieve additional space savings using this option. Where as,
by default the option `-scan-old-data` saves some time by skipping the
shared blocks.

``[-a, -shared-blocks <true>]`` - Compress Shared Blocks (if scanning old
data) [.advaced-priv]#(privilege: advanced)#::
Compresses the Compression Groups that have shared blocks created by
deduplication or cloning data. Valid only if `scan-old-data` parameter is
true.

``[-b, -snapshot-blocks <true>]`` - Compress Blocks In Snapshots (if
scanning old data) [.advaced-priv]#(privilege: advanced)#::
Compresses data blocks locked in a snapshot. Valid only if `scan-old-data`
parameter is true.

``[-q, -queue <true>]`` - Operation Should Be Queued::
Queues an efficiency operation. It will be queued only if an operation is
already in progress. Valid only if `scan-old-data` is false.

``[-f, -force <true>]`` - Force Operation::
Suppresses all confirmation messages.

``[-z, -skip-zero-replacement <true>]`` - Skip Zero block detection and
replacement [.advaced-priv]#(privilege: advanced)#::
Skip the zero block detection and replacement during the gatherer scan.
Valid only if `scan-old-data` parameter is true.

== Examples

The following examples start efficiency on a volume:
```

cluster1::> volume efficiency start -volume vol1 -vserver vs1

```

```

cluster1::> volume efficiency start -scan-old-data -volume vol1 -vserver vs1

```

```

cluster1::> volume efficiency start -volume vol1 -vserver vs1 -queue -delete-checkpoint

```

```
[[ID30362af9d08114ca9ecf3a921b1611f4]]
= volume efficiency stat
:hardbreaks:
```
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show volume efficiency statistics

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume efficiency stat` command displays efficiency statistics. The
output depends on the parameters specified with the command. If no
parameters are specified, the command displays the following efficiency
statistics fields for all the volumes:

* Vserver: The Vserver that the volume belongs to.
* Volume Name: Name of the volume.
* Inline Compression Attempts: Number of inline compression attempts done.
* Inline Incompressible CGs: Number of compression groups that cannot be
compressed by inline compression.



To display detailed information, run the command with `-instance`
parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed. The Vserver and
volume name are the default fields.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name::
Displays statistics only for those volume(s) that match the specified
Vserver.

{ ``[-volume <volume name>]`` - Volume Name::
Displays statistics only for those volume(s) that match the specified
volume name.

| ``[-path </vol/volume>]`` - Volume Path }::
```

```
Displays statistics only for those volume(s) that match the specified
volume path.

``[-b <true>]`` - Display In Blocks::
Displays usage size in 4k block counts.

``[-num-compressed-inline <integer>]`` - Inline Compression Attempts::
Displays statistics only for those volume(s) that match the specified
number of Compression Groups attempted inline.

``[-num-icg-inline <integer>]`` - Incompressible CGs Found By Inline
Compression::
Displays statistics only for those volume(s) that match the specified
number of incompressible Compression Groups encountered during Inline
Compression.

``[-inactive-blocks-compressed <integer>]`` - Number of Inactive Blocks
Compressed::
Displays statistics only for those volumes that match the specified number
of cold blocks on which compression was done.

== Examples

The following example displays default efficiency statistics for all the
volumes.
```

cluster1::> volume efficiency stat
Vserver: vs1
Volume: vol2
Inline Compression Attempts: 0
Inline Incompressible CGs: 0

Vserver: vs1
Volume: vol3
Inline Compression Attempts: 0
Inline Incompressible CGs: 0

```
[[ID4a1d50456a1cce799f1e33da7e9dac61]]
= volume efficiency stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop efficiency operation on a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `volume efficiency stop` command to stop an efficiency operation.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which efficiency operation needs
to be stopped.

| ``-path </vol/volume>`` - Volume Path }::
This specifies the volume path on which efficiency operation needs to be
stopped.

``[-a, -all <true>]`` - Stop All Operations::
This specifies both active and queued efficiency operations to be aborted.

== Examples

The following examples stop efficiency on a volume.
```

cluster1::> volume efficiency stop -vserver vs1 -volume vol1

```

```

cluster1::> volume efficiency stop -vserver vs1 -volume vol1 -all

```
[[ID96d776b60f7f97f2f2d2d5cbb0e7a5c4]]
= volume efficiency undo
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Undo efficiency on a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The command `volume efficiency undo` removes volume efficiency on a volume
by undoing compression, undoing compaction and removing all the block
sharing relationships, and cleaning up any volume efficiency specific data
structures. Any efficiency operations on the volume must be disabled
before issuing this command. The volume efficiency configuration is
deleted when the undo process completes. The command is used to revert a
volume to an earlier version of ONTAP where some of the efficiency
features are not supported. During this revert not all efficiencies needs
to be undone but only those gained by that particular feature (for
example, compaction), which is not supported in the earlier version.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume is located.

{ ``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume name.

| ``-path </vol/volume>`` - Volume Path [.advaced-priv]#(privilege:
advanced)# }::
This specifies the volume path.

``[-C, -compression <true>]`` - Decompress Data in the Volume [.advaced-
priv]#(privilege: advanced)#::
Undo the effects of compression. This requires efficiency to be disabled
(by performing xref:{relative_path}volume-efficiency-off.html[volume
efficiency off] ).

``[-D, -dedupe <true>]`` - Undo Block Sharing in the Volume [.advaced-
priv]#(privilege: advanced)#::
Undo the effects of deduplication. This requires efficiency to be disabled
```

```
(by performing xref:{relative_path}volume-efficiency-off.html[volume
efficiency off] ).

``[-i, -inode <integer>]`` - Inode Number to Undo Sharing [.advaced-
priv]#(privilege: advanced)#::
Remove the block sharings from a specified inode.

``[-t, -undo-type {all|wrong}]`` - Selective Undo [.advaced-
priv]#(privilege: advanced)#::
This specifies to remove either all or only invalid block sharing. When
``_all_`` is used, all block sharings are removed. When ``_wrong_`` is
used, only invalid sharings present in the volume are removed. When used
along with ``_log_`` option, it logs information about all or wrong block
sharings without sharing removal.

``[-d, -log <true>]`` - Only Log Incorrect Savings [.advaced-
priv]#(privilege: advanced)#::
If specified, information about invalid block sharing relationships will
only be logged. Invalid sharings will not be removed. This parameter is
only valid when the parameter `-undo-type` is specified as ``_wrong_`` .

``[-P, -data-compaction <true>]`` - Undo Data Compaction in the Volume
[.advaced-priv]#(privilege: advanced)#::
Undo the effects of data compaction.

``[-A, -cross-volume-dedupe <true>]`` - Undo Cross Volume Deduplication
[.advaced-priv]#(privilege: advanced)#::
Undo the effects of cross volume deduplication.

``[-X, -extended-compression <true>]`` - Extended Compression [.advaced-
priv]#(privilege: advanced)#::
Undo the effects of extended compression. This removes the compression
savings for data that requires more resources to compress.

``[-c, -auto-adaptive-compression <true>]`` - Auto Adaptive Compression
[.advaced-priv]#(privilege: advanced)#::
Undo the effects of auto adaptive compression.

``[-Z, -undo-compression-algorithm-list {lzopro|zstd|zlib}]`` - Undo
Compression Algorithm List [.advaced-priv]#(privilege: advanced)#::
Undo the effects of the specified compression algorithms.

``[-s, -extended-auto-adaptive-compression <true>]`` - Undo Extended Auto
Adaptive Compression [.advaced-priv]#(privilege: advanced)#::
Undo the effects of extended auto adaptive compression.

``[-a, -active-data-extended-auto-adaptive-compression <true>]`` - Undo
```

```
Active Data Extended Auto Adaptive Compression [.advaced-priv]#(privilege:
advanced)#::
Undo the effects of active data extended auto adaptive compression.

== Examples

The following are examples of how to use efficiency undo.

To undo deduplication and compression savings, but not compaction savings
in a volume name vol1 on a Vserver named vs1:
```

cluster1::*> volume efficiency undo -vserver vs1 -volume vol1

```
To rewrite compressed blocks and undo compression savings in a volume name
vol1 on a Vserver named vs1:
```

cluster1::*> volume efficiency undo -vserver vs1 -volume vol1 -compression

```
To rewrite compressed and deduped blocks without any efficiency in a
volume name vol1 on a Vserver named vs1:
```

cluster1::*> volume efficiency undo -vserver vs1 -volume vol1 -dedup -compression

```
To rewrite compacted blocks in a volume name vol1 on a Vserver named vs1:
```

cluster1::*> volume efficiency undo -vserver vs1 -volume vol1 -data-compaction

```
== Related Links

* xref:{relative_path}volume-efficiency-off.html[volume efficiency off]




[[ID863154d3619125c3cbc111543aff5522]]
= volume efficiency inactive-data-compression modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify volume inactive data compression configuration of a volume
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency inactive-data-compression modify` command is used
to modify the state of inactive data compression on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver of the volume.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume on which the inactive data compression state
should be modified.

``[-progress <text>]`` - Progress [.advaced-priv]#(privilege: advanced)#::
Progress of the scan.

``[-status <text>]`` - Status [.advaced-priv]#(privilege: advanced)#::
Operation on the scan succeeded or failed.

``[-failure-reason <text>]`` - Failure Reason [.advaced-priv]#(privilege:
advanced)#::
If scan failed, the reason of the failure.

``[-total-blocks <integer>]`` - Total Blocks to be Processed [.advaced-
priv]#(privilege: advanced)#::
Total number of blocks which must be processed by inactive data
compression scanner.

``[-total-processed <integer>]`` - Total Blocks Processed [.advaced-
priv]#(privilege: advanced)#::
Total number of blocks which are already processed by inactive data
compression scanner.

``[-percentage <percent>]`` - Progress [.advaced-priv]#(privilege:
advanced)#::
Scan progress in percentage.

``[-is-enabled {true|false}]`` - State of Inactive Data Compression on the
Volume [.advaced-priv]#(privilege: advanced)#::
Enable or disable inactive data compression on the volume. Disabling
```

```
inactive data compression is not allowed on Capacity optimized Flash with
QAT supported platforms.

``[-threshold-days <integer>]`` - Inactive data compression scan threshold
days value [.advaced-priv]#(privilege: advanced)#::
Threshold days value for inactive data compression scan. Valid values for
this option range from `-threshold-days-min` to `-threshold-days-max` .
This field is not supported on QAT supported platforms.

``[-threshold-days-min <integer>]`` - Inactive data compression scan
threshold minimum allowed value. [.advaced-priv]#(privilege: advanced)#::
Minimum allowed value in threshold days for inactive data compression
scan. Valid values for this option range from ``_1_`` to `-threshold-days-
max` . This field is not supported on QAT supported platforms.

``[-threshold-days-max <integer>]`` - Inactive data compression scan
threshold maximum allowed value. [.advaced-priv]#(privilege: advanced)#::
Maximum allowed value in threshold days for inactive data compression
scan. Valid values for this option range from `-threshold-days-min` to
``_60_`` . This field is not supported on QAT supported platforms.

``[-read-history-window-size <integer>]`` - Time window(in days) for which
client reads data is collected for tuning. [.advaced-priv]#(privilege:
advanced)#::
Client read history window. Valid values for this option range from
``_1_`` to ``_60_`` . This field is not supported on QAT supported
platforms.

``[-tuning-enabled {true|false}]`` - State of auto-tuning of Inactive data
compression scan on volume. [.advaced-priv]#(privilege: advanced)#::
Auto-tuning state of inactive data compression scan. This field is not
supported on QAT supported platforms.

``[-compression-algorithm {lzopro|zstd|zlib}]`` - Inactive data
compression algorithm [.advaced-priv]#(privilege: advanced)#::
Inactive data compression algorithm. The only supported algorithm is ZLIB
on QAT supported platforms.

== Examples

The following example displays information for modifying the state of
inactive data compression scan on volume "vol1":
```

cluster:::> volume efficiency inactive-data-compression modify -vserver vs1 -volume vol1 -is-enabled true
Inactive data compression scan enabled on volume vol1 in Vserver vs1.

```
[[ID46be5dd5d970a3a5227041a2a76ec721]]
= volume efficiency inactive-data-compression show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display volume inactive data compression progress

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency inactive-data-compression show` command is used to
show details of inactive data compression on a volume.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the vserver of the volume.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume on which inactive data compression should be
started.

``[-progress <text>]`` - Progress [.advaced-priv]#(privilege: advanced)#::
Progress of the scan.

``[-status <text>]`` - Status [.advaced-priv]#(privilege: advanced)#::
Operation on the scan succeeded or failed.

``[-failure-reason <text>]`` - Failure Reason [.advaced-priv]#(privilege:
advanced)#::
If scan failed, the reason of the failure.
```

```
``[-total-blocks <integer>]`` - Total Blocks to be Processed [.advaced-
priv]#(privilege: advanced)#::
Total number of blocks which must be processed by inactive data
compression scanner.

``[-total-processed <integer>]`` - Total Blocks Processed [.advaced-
priv]#(privilege: advanced)#::
Total number of blocks which are already processed by inactive data
compression scanner.

``[-percentage <percent>]`` - Progress [.advaced-priv]#(privilege:
advanced)#::
Scan progress in percentage.

``[-cold-blocks-found <integer>]`` - Number of Cold Blocks encountered
[.advaced-priv]#(privilege: advanced)#::
Number of cold blocks encountered by inactive data compression scanner.

``[-compression-done <integer>]`` - Number of Compression Done Blocks
[.advaced-priv]#(privilege: advanced)#::
Number of blocks on which compression was done.

``[-vol-overwrites <integer>]`` - Number of Vol-Overwrites [.advaced-
priv]#(privilege: advanced)#::
count of logical blocks overwritten in compressed data extents.

``[-is-enabled {true|false}]`` - State of Inactive Data Compression on the
Volume [.advaced-priv]#(privilege: advanced)#::
Is scan enabled on the volume.

``[-last-op-start-time <integer>]`` - Time since Last Inactive Data
Compression Scan started [.advaced-priv]#(privilege: advanced)#::
Time since the last inactive data compression scan op started.

``[-last-op-end-time <integer>]`` - Time since Last Inactive Data
Compression  Scan ended [.advaced-priv]#(privilege: advanced)#::
Time since the last inactive data compression scan op ended.

``[-last-succ-op-start-time <integer>]`` - Time since Last Successful
Inactive Data Compression Scan started [.advaced-priv]#(privilege:
advanced)#::
Time since the last successful inactive data compression scan op started.

``[-last-succ-op-end-time <integer>]`` - Time since Last Successful
Inactive Data Compression  Scan ended [.advaced-priv]#(privilege:
advanced)#::
```

```
Time since the last successful inactive data compression scan op ended.

``[-average-run-time <integer>]`` - Average run time for Inactive Data
Compression Scan [.advaced-priv]#(privilege: advanced)#::
Average runtime of inactive data compression.

``[-threshold-days <integer>]`` - Inactive data compression scan threshold
days value [.advaced-priv]#(privilege: advanced)#::
Threshold days value for inactive data compression scan. Valid values for
this option range from `-threshold-days-min` to `-threshold-days-max` .

``[-threshold-days-min <integer>]`` - Inactive data compression scan
threshold minimum allowed value. [.advaced-priv]#(privilege: advanced)#::
Minimum allowed value in threshold days for inactive data compression
scan. Valid values for this option range from ``_1_`` to `-threshold-days-
max` .

``[-threshold-days-max <integer>]`` - Inactive data compression scan
threshold maximum allowed value. [.advaced-priv]#(privilege: advanced)#::
Maximum allowed value in threshold days for inactive data compression
scan. Valid values for this option range from `-threshold-days-min` to
``_60_`` .

``[-read-history-window-size <integer>]`` - Time window(in days) for which
client reads data is collected for tuning. [.advaced-priv]#(privilege:
advanced)#::
Client read history window. Valid values for this option range from
``_1_`` to ``_60_`` .

``[-tuning-enabled {true|false}]`` - State of auto-tuning of Inactive data
compression scan on volume. [.advaced-priv]#(privilege: advanced)#::
Auto-tuning state of inactive data compression scan.

``[-compression-algorithm {lzopro|zstd|zlib}]`` - Inactive data
compression algorithm [.advaced-priv]#(privilege: advanced)#::
Inactive data compression algorithm.

``[-scan-mode
{default|compute_compression_savings|extended_recompression}]`` - Mode of
Inactive data compression scan [.advaced-priv]#(privilege: advanced)#::
Mode of the Inactive data compression scanner.

``[-phase1-l1s-processed <integer>]`` - Phase1 level1s Processed
[.advaced-priv]#(privilege: advanced)#::
Phase1 number of L1s processed by inactive data compression scanner.

``[-phase1-lns-skipped <text>,...]`` - Phase1 lns skipped [.advaced-
```

```
priv]#(privilege: advanced)#::
Phase1 number of indirect blocks skipped by inactive data compression
scanner.

``[-phase2-total-blocks <integer>]`` - Phase2 Total Blocks to be Processed
[.advaced-priv]#(privilege: advanced)#::
Phase2 number of total blocks which must be processed by inactive data
compression scanner.

``[-phase2-blocks-processed <integer>]`` - Phase2 Blocks Processed
[.advaced-priv]#(privilege: advanced)#::
Phase2 number of blocks processed by inactive data compression scanner.

``[-repacked-blocks <integer>]`` - Number of Repacked Blocks [.advaced-
priv]#(privilege: advanced)#::
Number of cold blocks repacked by inactive data compression scanner in
extended-recompression mode.

== Examples

The following example displays details about the inactive data compression
on volume "vol1":
```

cluster:::> volume efficiency inactive-data-compression show -vserver vs1 -volume vol1
Volume: vol1
Vserver: vs1
Is Enabled: true
Progress: IDLE
Status: SUCCESS
Compression Algorithm: lzopro
Failure Reason: -
Total Blocks: -
Total Blocks Processed: -
Phase1 L1s Processed: -
Phase1 Lns Skipped: -
Phase2 Total Blocks: -
Phase2 Blocks Processed: -
Percentage: -
Number of Cold Blocks Encountered: -
Number of Repacked Blocks: -
Number of Compression Done Blocks: -
Number of Vol-Overwrites: -
Time since Last Inactive Data Compression Scan Started(sec): -
Time since Last Inactive Data Compression Scan Ended(sec): -
Time since Last Successful Inactive Data Compression Scan Started(sec): -
Time since Last Successful Inactive Data Compression Scan Ended(sec): -
Average runtime of Inactive Data Compression(sec): -
Tuning Enabled: true
Threshold: 14
Threshold Upper Limit: 21

Threshold Lower Limit: 14
Client Read history window: 14

```
[[ID653118d616278b083b9baa55ac665b59]]
= volume efficiency inactive-data-compression start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start inactive data compression on a volume.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency inactive-data-compression start` command is used to
start inactive data compression on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the vserver of the volume.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume on which inactive data compression should be
started.

``[-s, -inactive-days <integer>]`` - Compression Days Count [.advaced-
priv]#(privilege: advanced)#::
Data older than 'inactive-days' days is compressed.

``[-m, -scan-mode
{default|compute_compression_savings|extended_recompression}]`` - scanner
mode [.advaced-priv]#(privilege: advanced)#::
This specifies in which mode inactive-data-compression-scanner should be
started. Three modes are available: 'default',
'compute_compression_savings' and 'extended_recompression'. 'default'
scanner will start compressing the cold data in volume.
'compute_compression_savings' scanner will calculate the auto adaptive
compression savings on the volume. 'extended_recompression' scanner will
attempt to re-write existing cold data to reduce internal fragmentation.
```

```
== Examples

The following example displays information for starting inactive data
compression scan on volume "vol1":
```

cluster:::> volume efficiency inactive-data-compression start -vserver vs1 -volume vol1
Inactive data compression scan started on volume vol1 in Vserver vs1.

```
[[ID09ffd9dc57c8c352d01cecfb2bc351ef]]
= volume efficiency inactive-data-compression stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop inactive data compression on a volume.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency inactive-data-compression stop` command is used to
stop inactive data compression on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the vserver of the volume.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume on which inactive data compression should be
stopped.

== Examples

The following example displays information for stopping inactive data
compression scan on volume "vol1":
```

cluster:::> volume efficiency inactive-data-compression stop -vserver vs1 -volume vol1
Inactive data compression scan stopped on volume vol1 in Vserver vs1.

```
[[IDdeced346547b0134540d9f281bda21a2]]
= volume efficiency inactive-data-reallocation modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Updates volume inactive data reallocation policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency inactive-data-reallocation modify` command is used
to modify the inactive data reallocation policy on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver of the volume.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume on which the inactive data compression state
should be modified.

``[-policy {none|default|on-read|all}]`` - Policy for Data Reallocation
[.advaced-priv]#(privilege: advanced)#::
This specifies policy for inactive data reallocation on the volume.
Inactive to active data reallocation is not supported on QAT supported
platforms.
+
Following are the possible values:
+
  * ``_default_`` : Reallocates data for random client driven data reads
on AFS.
  * ``_none_`` : Disables inactive to active data reallocation.
  * ``_on-read_`` : Reallocates data for client driven data reads
including sequential reads on AFS.



+
Currently unsupported policy:
+
```

```
  * ``_all_`` : Reallocates data for all client driven data reads on AFS
and snapshots. Also includes highly compressed data which was compressed
to one single 4k block.



== Examples

The following example displays information for modifying the inactive data
reallocation policy on volume "vol1":
```

cluster:::> volume efficiency inactive-data-reallocation modify -vserver vs1 -volume vol1 -policy default

```
[[ID49dcd9d07572c015701f230faddc7848]]
= volume efficiency inactive-data-reallocation show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display volume inactive data reallocation policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume efficiency inactive-data-reallocation show` command is used to
show details of inactive data reallocation policy on a volume.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the vserver of the volume.
```

```
``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume on which inactive data compression should be
started.

``[-policy {none|default|on-read|all}]`` - Policy for Data Reallocation
[.advaced-priv]#(privilege: advanced)#::
This specifies policy for inactive data reallocation on the volume.
+
Following are the possible values:
+
  * ``_default_`` : Reallocates data for random client driven data reads
on AFS.
  * ``_none_`` : Disables inactive to active data reallocation.
  * ``_on-read_`` : Reallocates data for client driven data reads
including sequential reads on AFS.


+
Currently unsupported policy:
+
  * ``_all_`` : Reallocates data for all client driven data reads on AFS
and snapshots. Also includes highly compressed data which was compressed
to one single 4k block.



== Examples

The following example displays details about the inactive data
reallocation on volume "vol1":
```

cluster:::> volume efficiency inactive-data-reallocation show -vserver vs1 -volume vol1
Volume: vol1
Vserver: vs1
Policy: default

```
[[ID793a63bfcf0095ba31d0500f181c5859]]
= volume efficiency policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an efficiency policy
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume efficiency policy create` creates an efficiency policy.

== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver on which the volume is located.

``-policy <text>`` - Efficiency Policy Name::
This specifies the policy name.

``[-type {threshold|scheduled}]`` - Policy Type::
This specifies the policy type. The policy type defines when the volume
using this policy will start processing a changelog. There are two
possible values:
+
  * ``_threshold_`` means changelog processing occurs when the changelog
reaches a certain percentage.
  * ``_scheduled_`` means changelog processing will be triggered by time.

+
The default value is ``_scheduled_`` .

``[-schedule <text>]`` - Job Schedule Name::
This specifies the job schedule. Use `job schedule` commands to manage job
schedules. Only cron job schedules are supported.

``[-duration <text>]`` - Duration::
This specifies the duration that an efficiency operation can run (in
hours). The possible values are "-" or a number between 1 and 999
inclusive. Default value is "-", which means no duration.

``[-start-threshold-percent <percent>]`` - Threshold Percentage::
The percentage at which the changelog will be processed. The percentage is
checked on an hourly basis. The default value is 20. Valid only if `-type`
parameter is set as ``_threshold_`` .

``[-qos-policy {background|best_effort}]`` - QoS Policy::
This specifies how the efficiency operations are throttled. This option
can be configured to be `background` or `best-effort` . Default value is
`best-effort` . If `background` is specified, the efficiency operations
are run with minimum or no impact on the data serving client operations.
```

```
If `best-effort` is specified, the efficiency operations might have some
impact on the data serving client operations.

``[-enabled {true|false}]`` - Enabled::
This specifies whether the policy is enabled or not. The policy is enabled
by default.

``[-comment <text>]`` - Comment::
User specified comment.

== Examples

The following example creates an efficiency policy.
```

cluster1::> volume efficiency policy create -vserver vs1 -policy policy1 -schedule daily -duration 100

```
[[IDc5f5c94347fc334a628b714743e71174]]
= volume efficiency policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an efficiency policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume efficiency policy delete` command deletes an efficiency
policy. An efficiency policy can be deleted only when it is not associated
with any volume. The pre-defined policies ``_default_`` and ``_inline-
only_`` cannot be deleted.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver on which the volume is located.

``-policy <text>`` - Efficiency Policy Name::
This specifies the policy name.

== Examples

The following example deletes an efficiency policy:
```

cluster1::> volume efficiency policy delete -vserver vs1 -policy policy1

```
[[ID4f99f1445a1d20757b4fa514532420c5]]
= volume efficiency policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an efficiency policy

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

The `volume efficiency policy modify` command can be used to modify the
policy attributes.

The attributes of the ``_inline-only_`` predefined policy cannot be
modified.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver on which the volume is located.

``-policy <text>`` - Efficiency Policy Name::
This specifies the policy name.

``[-type {threshold|scheduled}]`` - Policy Type::
This specifies the policy type. The policy type defines when the volume
using this policy will start processing a changelog. There are two
possible values:
+
  * ``_threshold_`` means changelog processing occurs when the changelog
reaches a certain percentage.
  * ``_scheduled_`` means changelog processing will be triggered by time.

+
The default value is ``_scheduled_`` .

``[-schedule <text>]`` - Job Schedule Name::
This specifies the job schedule. Use xref:{relative_path}job-schedule-
show.html[job schedule show] to show all the jobs.

``[-duration <text>]`` - Duration::
This specifies the duration that an efficiency operation can run in hours.
The possible value is between 1 and 999 inclusive.

``[-start-threshold-percent <percent>]`` - Threshold Percentage::
The percentage at which the changelog will be processed. The percentage is
checked on an hourly basis. The default value is 20. Valid only if `-type`
parameter is set as ``_threshold_`` .

``[-qos-policy {background|best_effort}]`` - QoS Policy::
This specifies how the efficiency operations are throttled. This option
can be configured to be `background` or `best-effort` . Default value is
`best-effort` . If `background` is specified, the efficiency operations
```

```
are run with minimum or no impact on the data serving client operations.
If `best-effort` is specified, the efficiency operations might have some
impact on the data serving client operations.

``[-enabled {true|false}]`` - Enabled::
This specifies whether the policy is enabled or not. Default value is
true.

``[-comment <text>]`` - Comment::
User specified comment.

== Examples

The following example modifies efficiency policy.
```

cluster1::> volume efficiency policy modify -policy policy1 -schedule hourly

```
== Related Links

* xref:{relative_path}job-schedule-show.html[job schedule show]




[[ID6d13b42e486272c6dbfffd2a99817617]]
= volume efficiency policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show efficiency policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume efficiency policy show` command displays information about
efficiency policies. By default, the command displays the following
information about all policies:

* Vserver: Name of the Vserver that the policy belongs to.
* Policy Name: Efficiency policy name.
* Job Schedule: Job schedule name.
```

```
* Duration (Hours): The duration in hours that the efficiency operation
can run.
* Enable: Whether the policy is enabled or not.
* Comment: User specified comment.



You can specify additional parameters to select the displayed information.
For example, to display efficiency policies only with duration 5 hours,
run the command with the `-duration 5` parameter.

The pre-defined policies ``_default_`` and ``_inline-only_`` are available
when all the nodes in the cluster are running ONTAP version 8.3 or later.

The ``_inline-only_`` pre-defined policy must be used when the user wants
to use the inline compression feature without any regularly scheduled or
manually started background storage efficiency operations. When a volume
is configured to use the ``_inline-only_`` efficiency policy, the system
will stop monitoring changes to the data for running the background
efficiency operations on the volume. Volumes cannot be configured with the
``_inline-only_`` policy if there is a currently active background
efficiency operation.

== Parameters

{ ``[-fields <fieldname>,...]``::
Selects the fields to be displayed. Vserver and policy are the default
fields (see example).

| ``[-instance  ]`` }::
If this parameter is specified, the command displays information about all
entries.

``[-vserver <vserver name>]`` - Vserver::
Selects information about the policies that match the specified Vserver.

``[-policy <text>]`` - Efficiency Policy Name::
Selects information about the policies that match the specified policy
name.

``[-type {threshold|scheduled}]`` - Policy Type::
Selects information about the policies that match the specified policy
type. There are two possible values - ``_threshold_`` and ``_scheduled_``
.

``[-schedule <text>]`` - Job Schedule Name::
Selects information about the policies that match the specified schedule.
```

```
``[-duration <text>]`` - Duration::
Selects information about the policies that match the specified duration
hours.

``[-start-threshold-percent <percent>]`` - Threshold Percentage::
Selects information about the policies that match the specified start-
threshold-percent. Valid only if `-type` parameter is set as
``_threshold_`` .

``[-qos-policy {background|best_effort}]`` - QoS Policy::
Selects information about the policies that match the specified throttling
method. The values can be `background` or `best-effort` .

``[-enabled {true|false}]`` - Enabled::
Selects information about the policies that have the specified enabled
setting.

``[-comment <text>]`` - Comment::
Selects information about the policies that match the specified comment.

``[-policy-owner {cluster-admin|vserver-admin}]`` - Owner of the Policy::
Selects information about the policies that match the specified owner. The
values can be `cluster-admin` or `vserver-admin` .

== Examples

The following example shows all the efficiency policies with the matching
Vserver vs1.
```

cluster1::> volume efficiency policy show -vserver vs1
Policy Job Duration
Vserver Name Schedule (Hours) QoS Policy Enabled Comment
-------- ------------ ---------- -------- ----------- -------- --------------
vs1 default daily - best_effort true Default policy
vs1 inline-only - - - - Inline-Only
policy
vs1 policy1 daily - best_effort true user-defined
3 entries were displayed.

```
The following example shows all the policies with the following fields -
Vserver (default), policy (default) and duration.
```

cluster1::> volume efficiency policy show -fields duration
vserver policy duration
------- ------- --------
vs1 default -

vs1 inline-only -
vs1 policy1 -
3 entries were displayed.

```
:leveloffset: -1


= volume encryption commands

:leveloffset: +1


[[ID45df05806640ad49f7aa48f193db315a]]
= volume encryption conversion pause
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Pause a running volume encryption conversion operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume encryption conversion pause` command pauses the running
encryption conversion operation on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume being encrypted.

``[-ignore-warning {true|false}]`` - Ignore Warning for Conversion Pause::
If this parameter is set, the command ignores the confirmation message.

== Examples
```

```
[[IDfba44fbde9e968309f71850c8fd20edf]]
= volume encryption conversion resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume a paused volume encryption conversion operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume encryption conversion resume` command resumes the paused
encryption conversion operation on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume being encrypted.

== Examples





[[IDcc9c8af44e61af3e910e7e293b9219f1]]
= volume encryption conversion show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show status of a volume encryption conversion

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `volume encryption conversion show` command displays information about
volume encryption conversion in the cluster. By default, with no
parameters, it only shows volume encryption operations that have failed or
are currently running. The command display output depends on the
parameters passed. If `-vserver` and `-volume` are specified, the
following information is displayed:

* Vserver Name: The Vserver on which the volume is located.
* Volume Name: The volume that is part of a completed or running volume
move operation.
* Start Time: The date and time when the volume encryption operation was
started.
* Status of Operation: The status of the operation.
* Percentage Completed: The amount of work to encrypt the volume completed
thus far in terms of percentage.



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``[-volume <volume name>]`` - Volume Name::
This parameter specifies the name of the volume being encrypted.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Start Time::
If this parameter is specified, the command displays encryption operations
that match the specified date and time in the cluster time zone when the
volume move operation started.

``[-status <text>]`` - Status::
If this parameter is specified, the command displays encryption operations
that match the specified status of the encryption operation.

== Examples
```

```
The following example shows a sample output for this command:
```

cluster1::> volume encryption conversion show

Vserver Volume Start Time Status
---------- ------------ --------------------- ----------------------
vs1 p2 9/18/2017 17:44:36 Phase 2 of 2 (redirect
scan) is in progress.

```
[[ID322063b040ce754bc91abd81af04b0b7]]
= volume encryption conversion start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a volume encryption conversion operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume encryption conversion start` command converts a non-encrypted
volume to encrypted volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume being encrypted.

``[-ignore-warning {true|false}]`` - Ignore Warning for Conversion Start::
If this parameter is set, the command ignores the confirmation message.

== Examples
```

```
[[ID42b590bd6ecc9e1c004e59b462f9bd6f]]
= volume encryption rekey pause
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Pause a running volume encryption rekey operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume encryption rekey pause` command pauses the running encryption
rekey operation on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume being encrypted.

``[-ignore-warning {true|false}]`` - Ignore Warning for Rekey Pause::
If this parameter is set, the command ignores the confirmation message.

== Examples





[[IDc0e32fafb67a36d721e9f89ad2c7dc03]]
= volume encryption rekey resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume a paused volume encryption rekey operation

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

The `volume encryption rekey resume` command resumes the paused encryption
rekey operation on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume being encrypted.

== Examples





[[ID11e38f0f633bb3c2c7491c55f938082e]]
= volume encryption rekey show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show status of a volume encryption rekey

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume encryption rekey show` command displays information about
volume encryption rekey in the cluster. By default, with no parameters, it
only shows volume encryption rekey operations that have failed or are
currently running. The command display output depends on the parameters
passed. If `-vserver` and `-volume` are specified, the following
information is displayed:

* Vserver Name: The Vserver on which the volume is located.
* Volume Name: The volume that is part of a completed or running volume
move operation.
```

```
* Start Time: The date and time when the volume encryption operation was
started.
* Status of Operation: The status of the operation.
* Percentage Completed: The amount of work to encrypt the volume completed
thus far in terms of percentage.



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``[-volume <volume name>]`` - Volume Name::
This parameter specifies the name of the volume being encrypted.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Start Time::
If this parameter is specified, the command displays encryption operations
that match the specified date and time in the cluster time zone when the
volume move operation started.

``[-status <text>]`` - Status::
If this parameter is specified, the command displays encryption operations
that match the specified status of the encryption operation.

== Examples

The following example shows a sample output for this command:
```

cluster1::> volume encryption rekey show

Vserver Volume Start Time Status
--------- ------------ -------------------- ----------------------
vs1 vol2 9/18/2017 17:51:41 Phase 2 of 2 (redirect
scan) is in progress.

```
[[IDc8cbaff5ab34f3422f00456dc1e69a13]]
```

```
= volume encryption rekey start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a volume encryption rekey operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume encryption rekey start` command changes the encryption key of
a volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume being rekeyed.

``[-ignore-warning {true|false}]`` - Ignore Warning for Rekey Start::
If this parameter is set, the command ignores the confirmation message.

== Examples





[[ID7198b040519c4ae89268c96dc954754b]]
= volume encryption secure-purge abort
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Abort secure deletion of trash in existing volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.
```

```
== Description

The `volume encryption secure-purge abort` command aborts the secure purge
operation on a volume.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the volume being encrypted.

== Examples





[[IDf1eb2626281523e126cff0253f5f9705]]
= volume encryption secure-purge show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show status of secure-purge operation on a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume encryption secure-purge show` command displays information
about volume encryption securepurge operation in the cluster.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver on which the volume is located.

``[-volume <volume name>]`` - Volume [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the volume being secure purged.

``[-status {invalid|initializing|snapshots-deleting|snapshots-
deleted|zombies-draining|zombies-drained|batched-free-log-
draining|batched-free-log-drained|finishing-trash-purge|finished-trash-
purge|reencrypting|aborting|aborted|success|failure|blocks-reclaim-scan-
in-progress|blocks-reclaim-scan-completed|updating-blocks-used}]`` -
Status [.advaced-priv]#(privilege: advanced)#::
This parameter displays the status of the secure purge operation.

== Examples

The following example shows a sample output for this command:
```

cluster1::> volume encryption secure-purge show
Vserver Volume Secure Purge Phase
---------- ------------ ----------
vs1 vol1 reencrypting

```
[[ID378623dfb580f2ca293245778cafc8e3]]
= volume encryption secure-purge start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start secure deletion of trash in existing volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume encryption secure-purge start` command performs secure purge
```

```
of encrypted volume.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the volume being encrypted.

{ ``[-delete-all-snapshots {true|false}]`` - Delete all Snapshot Copies?
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies if it is required to delete all the snapshots in
the volume.

``[-re-encryption-method {volume-move|in-place-rekey}]`` - Method to Use
for Re-keying the Volume [.advaced-priv]#(privilege: advanced)#::
Specifies the re-encryption method of the secure-purge operation. Possible
values are:
+
  *  volume-move
  *  in-place-rekey

+
The volume-move method is faster compared to in-place-rekey method. The
volume-move method requires additional space in the same aggregate. The
volume-move method will cause unforced MetroCluster switchovers to be
blocked while it is running. The volume-move operation involves cut-over
phase in which client I/O operations are temporary blocked for short
duration - less than 45 seconds. There is no cut-over phase for in-place-
rekey method.

| ``[-prepare <true>]`` - Method to Use for Preparing Secure Purge
[.advaced-priv]#(privilege: advanced)# }::
This parameter specifies if it is required to prepare secure purge of
encrypted volume.

== Examples





:leveloffset: -1
```

```
= volume file commands

:leveloffset: +1


[[ID374e430328cc4f892c539f3e909b1951]]
= volume file compact-data
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Apply Adaptive Data Compaction to a Snapshot copy of a file

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `volume file compact-data` command applies the Adaptive Data
Compaction feature to the snapshot of a file such that partially filled
blocks from that file will merge and consume less storage space.

== Parameters

``-node <nodename>`` - Node::
This parameter indicates the node name that the AWA instance runs on.

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the target file is located.

``-file </vol/<volume name>/<file path>>`` - File Path::
This specifies the complete file path. The snapshot name can be specified
as part of the path or by specifying the `-snapshot` parameter.

``[-volume <volume name>]`` - Volume  Name::
This specifies the volume in which the targeted file is located.

``[-snapshot <snapshot name>]`` - Snapshot Copy Name::
This specifies the snapshot name in which the file will be compacted.

== Examples

The following command applies the Adaptive Data Compaction feature to the
```

```
snapshot ``_snap1_`` of the file ``_/file1_`` in volume ``_vol1_`` :
```

cluster1::> volume file compact-data -vserver vs1 -volume vol1 -file /vol/vol1/file1 -snapshot snap1

```
[[IDfe205e4be333e5e849020bcf9e062584]]
= volume file modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Manage the association of a QoS policy group with a file

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command adds and removes files from QoS policy groups. QoS policy
groups define measurable service level objectives (SLOs) that apply to the
storage objects with which the policy group is associated. A QoS policy
group associated with this file can be created, modified, and deleted. You
cannot associate a file to a QoS policy group if a LUN was created from
the file.

== Parameters

``-vserver <vserver name>`` - Vserver Managing Volume::
This specifies the Vserver on which the volume (containing the file)
resides.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume. The name must be unique within the
hosting Vserver.

``-file <text>`` - File Path::
This specifies the actual path of the file with respect to the volume.

{ ``[-qos-policy-group <text>]`` - QoS Policy Group Name::
This option associates the file with a QoS policy group. This policy group
manages storage system resources to deliver your desired level of service.
If you do not assign a policy to a file, the system will not monitor and
control the traffic to it. To remove this file from a QoS policy group,
enter the reserved keyword “none”.
```

```
| ``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name
}::
This optional parameter specifies which QoS adaptive policy group to apply
to the file. This policy group defines measurable service level objectives
(SLOs) and Service Level Agreements (SLAs) that adjust based on the file's
allocated space or used space. To remove this file from an adaptive policy
group, enter the reserved keyword "none".

``[-caching-policy <text>]`` - Caching Policy Name::
 This optionally specifies the caching policy to apply to the file. A
caching policy defines how the system caches this volume's data in Flash
Cache modules. If a caching policy is not assigned to this file, the
system uses the caching policy that is assigned to the containing volume.
If a caching policy is not assigned to the containing volume, the system
uses the caching policy that is assigned to the containing Vserver. If a
caching policy is not assigned to the containing Vserver, the system uses
the default cluster-wide policy. The available caching policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read, and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.

+
Default caching-policy is auto.

== Examples
```

cluster1::> vol file modify -vserver vs0 -volume vs0_vol56 -file 1.txt -qos-policy-group fast -cache all-read

```
Associates the file ``_1.txt_`` with the ``_fast_`` QoS policy group and
``_all-read_`` caching policy.




[[ID021195b914124cd10f3b0248456ffe63]]
= volume file privileged-delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Perform a privileged-delete operation on unexpired WORM files on a
SnapLock enterprise volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The ` volume file privileged-delete ` command is used to perform a
privileged-delete operation on unexpired WORM files on a SnapLock
enterprise volume. The only built-in role that has access to the command
is ``_"vsadmin-snaplock"_`` .

== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver which hosts the SnapLock enterprise volume.

``-file </vol/<volume name>/<file path>>`` - File Path::
Specifies the absolute path of the file to be deleted. The value begins
with /vol/<volumename>.

== Examples

The following example deletes the unexpired WORM file
``_"/vol/vol1/wormfile"_`` . The file wormfile is stored in volume vol1
under Vserver ``_vserver1_`` .
```

vserver1::> volume file privileged-delete -file /vol/vol1/wormfile
[Job 76] Job succeeded: Privileged-delete of File "vs1:/vol/sle_vol1/wormfile" Completed.

```
[[ID26d1302779037300068bfc1d7aec8c9c]]
= volume file reservation
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Get/Set the space reservation info for the named file.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file reservation` command can be used to query the space
reservation settings for the named file, or to modify those settings. With
no further modifiers, the command will report the current setting of the
space reservation flag for a file. This tells whether or not space is
reserved to fill holes in the file and to overwrite existing portions of
the file that are also stored in a snapshot. For symlinks, the link is
followed and the command operates on the link target.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Specifies the Vserver on which the volume is located. If only one data
Vserver exists, you do not need to specify this parameter.

``-path </vol/<volume name>/<file path>>`` - File Name::
Specifies the complete file path for which we want to get/set the space
reservation settings.

``[-is-enabled <text>]`` - enable | disable::
Specifying enable or disable will turn the reservation setting on or off
accordingly for the file.

== Examples

The following example enables the file reservation setting for the file
named file1. The file file1 is stored in volume testvol on Vserver vs0.
```

node::> file reservation -vserver vs0 /vol/testvol/file1 enable
space reservations for file /vol/testvol/file1: on.

```
[[IDca0bcde8a9203928a0d5b67def892ddf]]
= volume file show-disk-usage
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show disk usage of file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command requires a path to a file in a volume and displays the
following information:

* Vserver name
* Total bytes used by the file in kilobytes
* Full Path to the file

If not logged in as Vserver administrator, the command also requires a
Vserver name.

NOTE: The "-instance" option provides the same result as the default as
there are no extra fields to display.



== 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.

| ``[-h  ]``::
If this parameter is specified, the command displays total bytes used by
the file in human readable form.

| ``[-k  ]``::
If this parameter is specified, the command displays total bytes used by
the file in kilobytes.

| ``[-m  ]``::
If this parameter is specified, the command displays total bytes used by
```

```
the file in megabytes.

| ``[-u  ]``::
If this parameter is specified, the command displays the unique bytes used
by the file (bytes that are not shared with any other file in the volume
due to deduplication or FlexClone files) in kilobytes.

| ``[-uh  ]``::
If this parameter is specified, the command displays the unique bytes used
by the file in human readable form.

| ``[-uk  ]``::
If this parameter is specified, the command displays the unique bytes used
by the file in kilobytes.

| ``[-um  ]``::
If this parameter is specified, the command displays the unique bytes used
by the file in megabytes.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``-vserver <vserver name>`` - Vserver::
This parameter is used to specify the Vserver that contains the file for
which the command displays the total bytes used. It is required if not
logged in as Vserver administrator.

``-path </vol/<volume name>/<file path>>`` - Full Path::
This required parameter is used to specify the path of the file for which
the command displays the total bytes used.

``[-r, -range <<start offset>:<end offset>>]`` - Block Range::
If this parameter is specified, the command displays the total bytes used
by the file in the specified block range.

== Examples

The following example displays the disk-usage of the file ``_file1.txt_``
in volume ``_/vol/root_vs0_`` .
```

cluster1::> volume file show-disk-usage -vserver vs0 -path /vol/root_vs0/file1.txt

Vserver Total Path
-------- ------ -----
vs0 1408KB /vol/root_vs0/file1.txt
cluster1::> volume file show-disk-usage -m -vserver vs0 -path /vol/root_vs0/file1.txt

Vserver Total Path
-------- ------ -----
vs0 1MB /vol/root_vs0/file1.txt
vs0::> volume file show-disk-usage -um -path /vol/root_vs0/file1.txt

Vserver Total Unique Path
-------- ----- ------ -------
vs0 1MB 1MB /vol/root_vs0/file1.txt

```
[[IDb0324584ae00dbe720162abc16bfacec]]
= volume file show-filehandle
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the file handle of a file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command requires a path to a file in a volume and displays the file
handle and volume information described below:

* Vserver name
* Path to the file
* File handle flags
* Snapshot ID of the file (snapid)
* File ID
* File handle generation number
* File system ID (fsid)
* Master data set ID (msid)
* Data set ID (dsid)
* Volume or constituent name

If not logged in as a Vserver administrator, the command also requires a
Vserver name.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Managing Volume [.advaced-
priv]#(privilege: advanced)#::
This specifies the Vserver where the file resides.

``[-path <text>]`` - Path to File [.advaced-priv]#(privilege: advanced)#::
This specifies the path to the file.

== Examples

The following example displays the file handle and volume information of a
file named ``_file1.txt_`` in the volume ``_/vol/vol1_`` .
```

cluster1::> volume file show-filehandle -vserver vs0 -path /vol/vol1/file1.txt
Vserver Path
--- ---
vs0 /vol/vol1/file1.txt
flags snapid fileid generation fsid msid dsid
--- --- --- --- --- --- ---
0x0 0 0x60 0x206b6 0x402 0x80000402 0x402
volume
---
vol1__0003

```
[[IDd7336a74636c820d15359e3537730519]]
= volume file show-inode
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display file paths for a given inode

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays information about all the files having a given inode
in a volume of a Vserver. If the `-snapshot-id` or `-snapshot-name`
parameter is specified, the command displays file information from the
snapshot; otherwise, it displays the information from the active file
```

```
system. The `-vserver` , `-volume` and `-inode-number` are mandatory
parameters.

If no optional parameter is specified, the command displays the following
fields for all the files having the given inode:

* Vserver Name
* Volume Name
* Inode Number
* File Path

The `volume file show-inode` command is only supported on flexible volumes
and FlexGroup constituents.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields parameter, the command output also includes the
specified field or fields.

| ``[-snapshot  ]``::
If this parameter is specified, the command displays the following
information:
+
  * Vserver Name
  * Volume Name
  * Inode Number
  * Snapshot Name
  * Snapshot ID
  * File Path



| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed information
about the files matching the specified inode number. The following
information is displayed:
+
  * Vserver Name
  * Volume Name
  * Inode Number
  * File Path
  * Snapshot Name
  * Snapshot ID
  * File Name
  * Parent Inode Number
  * Parent Directory Cookie
```

```
``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the volume or snapshot is located.

``-volume <volume name>`` - Volume Name::
This specifies the volume in which the inode number is located.

``-inode-number <integer>`` - Inode Number::
This specifies the inode number whose information has to be retrieved.

{ ``[-snapshot-name <snapshot name>]`` - Snapshot Name::
If this parameter or `-snapshot-id` is specified, information about the
files is retrieved from the snapshot instead of the active file system.

| ``[-snapshot-id <integer>]`` - Physical Snapshot ID }::
If this parameter or `-snapshot-name` is specified, information about the
files is retrieved from the snapshot instead of the active file system.

``[-file-path <text>]`` - File Path::
If this parameter is specified, the command displays information only
about the files that match the specified file path.

``[-file-name <text>]`` - File Name::
If this parameter is specified, the command displays information only
about the files that match the specified file name.

``[-parent-inode-number <integer>]`` - Parent Inode Number::
The inode number of the parent directory of the file associated with the
inode. If this parameter is specified, the command displays information
only about the files that match the specified parent inode number.

``[-parent-dir-cookie <integer>]`` - Parent Directory Cookie::
The index of the directory entry of the file in its parent directory tree.
If this parameter is specified, the command displays information only
about the files that match the specified parent directory cookie.

== Examples

The following example displays information about all the files having the
inode number 96 in the active file system of a volume named vol1 on a
Vserver named vs1:
```

cluster1::> volume file show-inode -vserver vs1 -volume vol1 -inode-number 96
Inode

Vserver Volume Number File Path
--------- ------------ -------- ----------
vs1 vol1 96 /vol/vol1/file1
vs1 vol1 96 /vol/vol1/file2
vs1 vol1 96 /vol/vol1/A/file2
3 entries were displayed.

```
The following example displays information about all the files with inode
number 96 in a snapshot named mysnap. The snapshot is present in a volume
named vol1 on a Vserver named vs1:
```

cluster1::> volume file show-inode -vserver vs1 -volume vol1 -inode-number 96 -snapshot-name mysnap -snapshot
Inode Snapshot Snapshot
Vserver Volume Number Name ID File Path
--------- ------------ -------- -------- -------- ----------
vs1 vol1 96 mysnap 131 /vol/vol1/.snapshot/mysnap/file1
vs1 vol1 96 mysnap 131 /vol/vol1/.snapshot/mysnap/file2
2 entries were displayed.

```
The following example displays detailed information about all the files
with inode number 96 in a snapshot named mysnap. The snapshot is present
in a volume named vol1 on a Vserver named vs1:
```

cluster1::> volume file show-inode -vserver vs1 -volume vol1 -inode-number 96 -snapshot-name mysnap -instance
Vserver Name: vs1
Volume Name: vol1
Inode number: 96
File Path: /vol/vol1/.snapshot/mysnap/file1
Snapshot Name: mysnap
Physical Snapshot ID: 131
File Name: file1
Parent Inode Number: 64
Parent Directory Cookie: 2
Vserver Name: vs1
Volume Name: vol1
Inode number: 96
File Path: /vol/vol1/.snapshot/mysnap/file2
Snapshot Name: mysnap
Physical Snapshot ID: 131
File Name: file2
Parent Inode Number: 64
Parent Directory Cookie: 3
2 entries were displayed.

```
[[IDa88960fefd8096443210abc395bf0518]]
= volume file async-delete cancel
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cancel an async directory delete job

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume file async-delete cancel` command cancels an async delete of a
directory.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume containing the directory
resides.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume.

``-jobid <text>`` - Job ID [.advaced-priv]#(privilege: advanced)#::
This specifies the Job ID of the async directory delete job to cancel. Run
xref:{relative_path}volume-file-async-delete-show.html[volume file async-
delete show] to get the job ID.

== Examples

The following example cancels a job with job ID "0:1" in volume "vol1" in
Vserver "vs1".
```

cluster1::*> volume file async-delete cancel -vserver vs1 -volume vol1 -jobid 0:1

```
== Related Links

* xref:{relative_path}volume-file-async-delete-show.html[volume file
async-delete show]




[[ID056d1f6e71ce3c61ff47adbbdbb5e77c]]
= volume file async-delete prepare-for-revert
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cancel all async directory delete jobs before starting revert

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume file async-delete prepare_for_revert` command cancels all
async directory delete jobs.

== Examples

The following example deletes all async directory delete jobs.
```

cluster1::*> volume file async-delete prepare-for-revert

```
[[ID9e042861b4192b3e26da4b350553c32b]]
= volume file async-delete show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
List ongoing async directory delete jobs

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.
```

```
== Description

The event log will contain events for async directory delete operations
that have already completed. Use the ``_event log show -event
async_delete_`` command to find completed async directory delete
operations.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, ongoing async directory delete operations
that match the specified Vserver are displayed.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, ongoing async directory delete operations
that match the specified volume are displayed.

``[-jobids <text>,...]`` - Job ID [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, ongoing async directory delete operations
that match the specified job ID are displayed.

``[-paths <text>,...]`` - Directory Path [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, ongoing async directory delete operations
that match the specified path are displayed.

``[-deleted-files <integer>,...]`` - Number of files deleted [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the number of files deleted as part of the
corresponding async directory delete job in a FlexVol.

``[-deleted-dirs <integer>,...]`` - Number of directories deleted
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the number of directories deleted as part of the
corresponding async directory delete job in a FlexVol.
```

```
``[-deleted-bytes <integer>,...]`` - Deleted size in bytes [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies total number of bytes deleted as part of the
corresponding async directory delete job in a FlexVol.

== Examples

The following example shows that there are two async directory deletes in
progress; one in FlexVol `vol1` and another one in FlexGroup `fg` .
```

cluster1::*> volume file async-delete show
Vserver Volume Job ID Path Progress
(#File #Dir Size(B))
--- --- --- --- ---
vs1 vol1 0:2 d2/dd 1250 2 24560
vs2 fg 0:1 d1/d2 - - -

```
[[ID4898ed4268772af7ad2ce1c288a05d21]]
= volume file async-delete start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start async delete of a directory

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume file async-delete start` command starts async delete of a
directory.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume containing the directory
resides.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume.

``-path <text>`` - Directory Path [.advaced-priv]#(privilege: advanced)#::
This specifies the path of the directory.

``[-throttle <integer>]`` - Throttle [.advaced-priv]#(privilege:
advanced)#::
 This specifies maximum number of directory delete operations per second.

== Examples

The following example deletes directory "d2" under directory "d1" in
volume "vol1" hosted in Vserver "vs1".
```

cluster1::*> volume file async-delete start -vserver vs1 -volume vol1 -path d1/d2

```
[[IDae8fcee4cd87f5a0f6c6e0f257b3d613]]
= volume file async-delete client disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable async delete of a directory from the client

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume file async-delete client disable` command disables async
delete of a directory on the volume, from client applications.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume resides.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume.

== Examples

The following example disables async directory delete from the client in
volume "vol1" hosted in Vserver "vs1".
```

cluster1::*> volume file async-delete client disable -vserver vs1 -volume vol1

```
[[ID18d6c93bc6fd76512ada9e8fa6069d52]]
= volume file async-delete client enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable async delete of a directory from the client

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume file async-delete client enable` command enables async delete
of a directory on the volume, from client applications. When async
directory delete from the client is enabled, 'mv' or 'rename' of a
directory, to the trashbin name, triggers async delete of the directory.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume resides.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume.

``[-trashbin <text>]`` - Trashbin Directory Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the name of the trashbin directory.

== Examples

The following example enables async delete of directory in volume "vol1"
hosted in Vserver "vs1" from client.
```

cluster1::*> volume file async-delete client enable -vserver vs1 -volume vol1 -trashbin ntaptrash

```
[[IDb36ac5da64a707946414af0d19b10504]]
= volume file async-delete client show
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the status of async delete of a directory from the client

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume file async-delete client show` command shows the configuration
of async delete of a directory on a volume from client applications.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume resides.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the volume.

``[-is-enabled {true|false}]`` - Is Client Async Delete Enabled [.advaced-
priv]#(privilege: advanced)#::
This specifies whether async delete of a directory from client
applications is enabled on the volume.

``[-trashbin <text>]`` - Trashbin Directory Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the name of the trashbin directory.

``[-inode <integer>]`` - Trashbin Directory Inode Number [.advaced-
priv]#(privilege: advanced)#::
This specifies the inode number of the trashbin directory.

== Examples
```

```
The following example shows the enable status of async directory delete
from the client for volume "vol1" hosted in Vserver "vs1".
```

cluster1::*> volume file async-delete client show -vserver vs1 -volume vol1

```
[[ID6cab79849e7e28fb84ce0c56c0b9dc66]]
= volume file clone autodelete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable/Disable autodelete

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone autodelete` command enables or disables the
automatic deletion of file, LUN or NVMe namespace clones. Newly created
file, LUN and NVMe namespace clones are disabled for automatic deletion by
default.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume resides. If only one data
Vserver exists, you do not need to specify this parameter.

``[-volume <volume name>]`` - Volume Name::
This specifies the name of the volume in which the file, LUN or NVMe
namespace is present.

``-clone-path <text>`` - Clone Path::
This specifies the path where clone resides. If you use the volume
parameter, then specify the relative path to the file, LUN or NVMe
namespace clone. Otherwise, specify the absolute path.

``-enable {true|false}`` - Enable or Disable Autodelete::
This parameter enables or disables the autodelete feature for the file,
LUN or NVMe namespace clones in a specified volume if the clones are
already added for automatic deletion. If you set the parameter to true,
the specified file, LUN or NVMe namespace clones get automatically deleted
```

```
in the 'try' or 'disrupt' mode. If the value is false, the clones get
automatically deleted only in the 'destroy' mode.

``[-force <true>]`` - Force Enable or Disable Autodelete::
If `-enable` is ``_true_`` then this parameter forces automatic deletion
of a specified file, LUN or NVMe namespace, or a file, LUN or NVMe
namespace clone. If `-enable` is ``_false_`` then specifying this
parameter disables autodeletion on a file, LUN or NVMe namespace - or a
file, LUN or NVMe namespace clone - even if `-commitment```_destroy_`` is
specified.

== Examples

The following command enables for automatic deletion a LUN Clone named
``_lun_clone_`` contained in a volume named ``_volume1_`` . This volume is
present on a Vserver named ``_vs1_`` .
```

cluster1::> volume file clone autodelete /vol/volume1/lun_clone -enable true -vserver vs1

```
The following command specifies the relative clone path when the volume
parameter is specified in the command.
```

cluster1::> volume file clone autodelete lun_clone -enable true
-vserver vs1 -volume volume1

```
[[ID05fa140d8a81e8d151d498080ccbde2e]]
= volume file clone create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create file or LUN full or sub file clone

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone create` command creates a clone of a file, a LUN or
an NVMe namespace. You can optionally specify the following parameters for
the clone file creation process:

* Vserver in which the volume resides
```

```
* Name of the parent snapshot
* The range of blocks to be cloned
* The option to avoid space reservations for the new file or LUN clone
* The option to assign a QoS policy group to the new file or LUN clone
* The option to assign a caching policy to the new file or LUN clone
* The option to mark the new file, LUN or NVMe namespace clone created for
auto deletion
* The option to overwrite an existing file, LUN or NVMe namespace clone

File, LUN or NVMe namespace clones create a duplicate copy of another
file, LUN or NVMe namespace, but don't require copying the data itself.
This allows the clone operation to occur in constant time, taking the same
amount of time to complete no matter the size of the file being cloned.
This also means that clones require only a small amount of additional
storage space because the clone shares the data with the source file, LUN
or NVMe namespace.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver in which the parent volume resides. If only one
data Vserver exists, you do not need to specify this parameter.

``[-volume <volume name>]`` - Volume::
This specifies the name of volume in which a file, LUN or NVMe namespace
is going to be cloned.

``-source-path <text>`` - Source Path::
This specifies the path to the file, LUN or NVMe namespace to be cloned
relative to the specified volume.

``-destination-path <text>`` - Destination Path::
This specifies the path for the newly-created cloned file, LUN or NVMe
namespace relative to the specified volume. If the file, LUN or NVMe
namespace clone to be created is a whole file, LUN or NVMe namespace, the
destination file, LUN or NVMe namespace must not exist. If the `range`
parameter is specified, the destination file or LUN must exist. If the
`snapshot-name` parameter is specified, this option is mandatory.

``[-s, -snapshot-name <snapshot name>]`` - Source Snapshot::
The name of the snapshot to use as the source for the clone operation. If
this value is not specified, the active filesystem will be used instead.

{ ``[-r, -range <text>,...]`` - Block Range::
This specifies the block range to be cloned. If the range is not
specified, the entire file, LUN, or NVMe namespace is cloned. The block
```

```
range should be specified in the format s:d:n where s is the source start
block number, d is the destination start block number, and n is the length
in blocks to be cloned. The range of n should be from 1 to 32768 or 1 to
16777216 in case of clone from Active File System or snapshot
respectively. If this parameter is used in the path provided by the
`destination-path` , the parameter must refer to a file, LUN, or NVMe
namespace which already exists. If either the source or destination is a
LUN or NVMe namespace, then the block size is measured in LBA blocks. The
source object block size and destination block size must be equal. If
neither the source nor destination is a LUN or NVMe namespace, then the
block size will be 4KB. If 512-byte sectors are used, the source and
destination offsets must have the same offset within 4KB blocks.  + This
option is most likely to be used by external automated systems in managing
virtual disk configurations and not by human administrators.

| ``[-o, -no-reserve {true|false}]`` - Do not reserve clone }::
If this option is used, the clone file or LUN will not be guaranteed space
in the underlying aggregate. While this out-of-space condition persists,
writes to the clone file or LUN would fail. This option may be useful if
few writes to the clone are expected to be needed, or to allow a file or
LUN clone to be created under space-constrained conditions for recovery
purposes. If this option is not specified the clone will inherit the space
reservation properties from the source.

``[-i, -ignore-streams {true|false}]`` - Ignore streams::
This parameter specifies whether streams should be ignored during cloning
of files, LUNs, or NVMe namespaces. If you set this parameter to false,
the streams are ignored; otherwise, they are included in the clones. The
default value is false.

``[-k, -ignore-locks {true|false}]`` - Ignore locks::
This parameter specifies whether byte-range locks and shared-mode locks on
files, LUNs or NVMe namespaces should be ignored during cloning. If you
set this parameter to true, the locks are ignored; otherwise, clone
operation fails if locks are present on files, LUNs or NVMe namespaces.
The default value is false.

``[-d, -overwrite-destination {true|false}]`` - Overwrite Destination::
Specify this parameter to overwrite the destination file, LUN, or NVMe
namespace, if it exists. The default is not to overwrite the destination
file. The command will fail if the destination file exists.

{ ``[-qos-policy-group <text>]`` - QoS Policy Group Name::
 This optionally specifies which QoS policy group to apply to the file or
LUN. This policy group defines measurable service level objectives (SLOs)
that apply to the storage objects with which the policy group is
```

```
associated. If you do not assign a policy group to a file or LUN, the
system will not monitor and control the traffic to it. You cannot
associate a file to a QoS policy group if a LUN was created from the file.

| ``[-qos-adaptive-policy-group <text>]`` - QoS Adaptive Policy Group Name
}::
This optionally specifies which QoS adaptive policy group to apply to the
file or LUN. This policy group defines measurable service level objectives
(SLOs) and Service Level Agreements (SLAs) that adjust based on the file
or LUN's allocated space or used space.

``[-caching-policy <text>]`` - Caching Policy Name::
 This optionally specifies the caching policy to apply to the file. A
caching policy defines how the system caches this volume's data in Flash
Cache modules. If a caching policy is not assigned to this file, the
system uses the caching policy that is assigned to the containing volume.
If a caching policy is not assigned to the containing volume, the system
uses the caching policy that is assigned to the containing Vserver. If a
caching policy is not assigned to the containing Vserver, the system uses
the default cluster-wide policy. The available caching policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read, and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.

+
Default caching-policy is auto.

``[-autodelete {true|false}]`` - Mark Clone for Autodeletion::
This parameter marks the file, LUN or NVMe namespace clones created for
auto deletion. When set to true, the file, LUN or NVMe namespace clones
get automatically deleted when the volume runs out of space. The default
value is false.

``[-bypass-throttle {true|false}]`` - Bypass Throttle Checks [.advaced-
```

```
priv]#(privilege: advanced)#::
This parameter specifies whether clone throttle checks should be skipped
during clone creation. When set to true, clones are created without
enforcing any clone throttle checks. The default value is false.

``[-nosplit-entry {true|false}]`` - Mark Clone to keep Unsplit::
This parameter specifies whether the clone should be kept unsplit. When
set to true, clones created will not undergo split routine. It is
applicable only for full-file clones created from Active File System. The
default value is false.

``[-is-backup {true|false}]`` - Is a Clone for Backup::
This parameter is used to mark the destination file as a backup clone,
where divergence is expected on the source file and no divergence is
expected on the destination file. It is applicable only for full-file
clones created from Active File System volumes. The default value is
``_false_`` .

``[-destination-volume <volume name>]`` - Destination Volume::
This specifies the name of the volume where the destination file resides.
This can be different from `volume` , whereas parameter `volume` specifies
the volume on which source file resides. This is an optional argument that
applies only to a MetaWAFL volume where the source and destination volumes
for the clone operation can be different. If this parameter is not given,
the destination file will be created in the volume where source_file
resides.

== Examples

The following command creates a FlexClone file of the file named
``_myfile_`` contained in a volume named ``_vol_`` . The file ``_myfile_``
is located in the root directory of that volume. The cloned file
``_myfile_copy_`` resides in the root directory same volume.
```

cluster1::> volume file clone create -volume vol -source-path /myfile -destination-path /myfile_copy

```
The following command optionally associates the FlexClone file named
``_myfile_copy_`` with the ``_fast_`` QoS policy group and the caching
policy named random-read.
```

cluster1::> volume file clone create -volume vol -source-path /myfile -destination-path /myfile_copy -qos-policy -group fast -caching-policy random-read

```
[[IDddfd31abd13cc332e8c8734d7691947b]]
= volume file clone show-autodelete
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the autodelete status for a file or LUN clone

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone show-autodelete` command displays the autodelete
details of a file, LUN or NVMe namespace clone. The command displays the
following information about a file, LUN or NVMe namespace clone:

* Vserver Name
* Clone Path
* Whether auto deletion of file, LUN or NVMe namespace clone is enabled



== 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.

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver to which the file, LUN or NVMe namespace clone
belongs.

``-clone-path <text>`` - Clone Path::
This specifies the path of the file, LUN or NVMe namespace clone.

``[-autodelete-enabled {true|false}]`` - Autodelete Enabled::
If this parameter is true, the file, LUN or NVMe namespace clone gets
automatically deleted in the 'try' or 'disrupt' mode. If the value is
false, the clones get automatically deleted only in the 'destroy' mode.

== Examples
```

```
The following example displays the autodelete information about a file,
LUN or NVMe namespace clone.
```

cluster1::> volume file clone show-autodelete -vserver vs1 -clone-path /vol/v1/f1
Vserver Name: vs1
Clone Path: /vol/v1/f1
Autodelete Enabled: true

```
[[IDd706d954ce57f368f5d641409450b1aa]]
= volume file clone deletion add-extension
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add new supported file extensions to be deleted with clone delete

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone deletion add-extension` command can be used to add
new supported file extensions for clone delete.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Name of the vserver.

``-volume <volume name>`` - Volume Name::
Name of the volume.

``-extensions <text>`` - Supported Extensions for Clone Delete::
List of supported file extensions for clone delete.

== Examples

The following example adds the new supported ``_vmdk,vhd_`` file
extensions to volume ``_vol1_`` of vserver ``_vs1_`` .
```

cluster1::> volume file clone deletion add-extension -vserver vs1 -volume vol1 -extensions vmdk,vhd

```
[[ID0b41d90ac31645bbb09b6e466ed1828d]]
= volume file clone deletion modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]


*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone deletion modify` command can be used to change the
required minimum clone file size of a volume for clone delete.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Name of the vserver.

``-volume <volume name>`` - Volume Name::
Name of the volume.

``[-minimum-size {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Size Required
for Clone delete::
Minimum clone file size required for clone delete.

== Examples

The following example changes the required minimum file size to 100M for
volume ``_vol1_`` of vserver ``_ vs1 _`` .
```

cluster1::> volume file clone deletion modify -volume vol1 -vserver vs1 -minimum-size 100M

```
[[ID4640b7818ecce12124c523e19ba1994d]]
= volume file clone deletion remove-extension
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove unsupported file extensions for clone delete

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone deletion remove-extension ` command can be used to
remove the existing file extensions that are no longer supported for clone
delete.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Name of the vserver.

``-volume <volume name>`` - Volume Name::
Name of the volume.

``[-extensions <text>]`` - Unsupported Extensions for Clone Delete::
List of unsupported file extensions for clone delete.

== Examples

The following example removes the existing unsupported ``_vmdk, vhd_``
file extensions to volume ``_vol1_`` of vserver ``_vs1_`` .
```

cluster1::> volume file clone deletion remove-extension -vserver vs1 -volume vol1 -extensions vmdk,vhd

```
[[IDe09d34c800ddcaec997203f864036c04]]
= volume file clone deletion show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Show the supported file extensions for clone delete

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone deletion show` command displays the following
information for clone delete:

* Vserver Name
* Volume Name
* Minimum File Size Required for Clone Delete
* List of Supported File Extensions for Clone Delete



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
Name of the vserver.

``[-volume <volume name>]`` - Volume Name::
Name of the volume.

``[-extensions <text>,...]`` - Supported Extensions for Clone Delete::
List of supported file extensions for Clone Delete.

``[-minimum-size {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum Size Required
for Clone delete::
Minimum file size required for Clone Delete.

== Examples

The following example displays the clone deletion information for all
volumes of all vservers.
```

cluster1::> volume file clone deletion show
Vserver Volume Minimum Extensions
Size
---------- ---------- ------- ----------
vs0 testvol 100B vmdk, vhd, vhdx, vswp
vs0_root 0B -
vs1 testvol 100G vmdk, vhd, vhdx, vswp
vs1_root 0B -

```
The following example displays the clone deletion information for volume
``_vol1_`` of vserver ``_vs1_`` .
```

cluster1::> volume file clone deletion show -vserver vs0 -volume testvol
Vserver Name: vs0
Volume Name: testvol
Supported Extensions for Clone Delete: vmdk, vhd, vhdx, vswp
Minimum Size Required for Clone delete: 100B

```
[[IDc6c4690dfe4be736b275841f60b5fbd2]]
= volume file clone split load modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify maximum split load on a node

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume file clone split load modify` command can be used to change
the maximum split load (file, LUN or NVMe namespace clones) of a node.

== Parameters

``-node {<nodename>|local}`` - Node Name::
Node name on which the new maximum split load is being applied.

``[-max-split-load {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Clone Split
Load::
This specifies the new maximum split load of a node. This is the amount of
clone create load, the node can take at any point of time. If it crosses
this limit, then the clone create requests will not be allowed, till the
split load is less than maximum split load

== Examples

The following example changes the new maximum limit to 10TB on node1.
```

cluster1::*> volume file clone split load*> modify -node clone-01 -max-split-load 100KB

```
[[ID7cd49f4284bd7ee5902534f5acb285de]]
= volume file clone split load show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show split load on a node
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file clone split load show` command displays the corresponding
file, LUN or NVMe namespace clone split loads on nodes.If no parameters
are specified, the command displays the following information:

* Node
* Max Split Load
* Current Split Load
* Token Reserved Load
* Allowable Split Load



== 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>|local}]`` - Node Name::
Node on which the file, LUN or NVMe namespace clone split load is
displayed.

``[-max-split-load {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Clone Split
Load::
This specifies the maximum allowable split load on the node.

``[-current-split-load {<integer>[KB|MB|GB|TB|PB]}]`` - Current Clone
Split Load::
This specifies the current on going split load on the node.

``[-token-reserved-load {<integer>[KB|MB|GB|TB|PB]}]`` - Load Reserved for
Clone Creation::
This specifies the reserved split load of the node using the tokens.

``[-allowable-split-load {<integer>[KB|MB|GB|TB|PB]}]`` - Allowable Clone
Split Load::
```

```
This specifies the available split load of the node.

== Examples

The following example displays the current and allowable file, LUN or NVMe
namespace clone split load on a node.
```

cluster1::> volume file clone split load show
Node Max Current Token Allowable
Split Load Split Load Reserved Load Split Load
---------- ---------- ---------- ---------- ----------
clone-01 15.97TB 0B 100MB 15.97TB
clone-02 15.97TB 0B 100MB 15.97TB
2 entries were displayed.
cluster1::> volume file clone split load show -node clone-01 -instance
Node Name: clone-01
Maximum Clone Split Load: 15.97TB
Current Clone Split Load: 0B
Load Reserved for Clone Creation: 100MB
Allowable Clone Split Load: 15.97TB

```
[[IDfaeb326186be1cc4870914f65300fdb2]]
= volume file fingerprint abort
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Abort a file fingerprint operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file fingerprint abort` command aborts an in-progress
fingerprint operation. This command only aborts the fingerprint operations
that have not yet completed. This command takes session-id as input and
aborts the fingerprint operation that is associated with that particular
session-id.

== Parameters

``-session-id <integer>`` - Session ID of Fingerprint Operation::
Specifies the session-id of the fingerprint operation that needs to be
aborted. It is an unique identifier for the fingerprint operation. This
session-id is returned when the fingerprint operation is started on a
file.

== Examples

The following example aborts the fingerprint operation identified by
``_17039361_`` :
```

cluster1::> volume file fingerprint abort -session-id 17039361

```
[[IDf8f09aa7f179d88c314d908ff214b3b9]]
= volume file fingerprint dump
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display fingerprint of a file
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file fingerprint dump` command displays the following
information given the `-session-id` of the fingerprint operation:

* Vserver:

The Vserver on which the file exists.
* Session-ID:

A unique identifier for the fingerprint operation. This session-id is
returned when the fingerprint operation is started on a file. The session-
id of the fingerprint operation can be used to get the progress of an
ongoing fingerprint operation as well as the complete fingerprint output
for the file once the operation is completed.
* Volume:

The name of the volume on which the file resides.
* Path:

The absolute path of the file on which the fingerprint is calculated. The
value begins with /vol/<volumename>.
* Data Fingerprint:

The digest value of data of the file. The fingerprint is base64 encoded.
This field is not included if the scope is ``_metadata-only_`` .
* Metadata Fingerprint:

The digest value of metadata of the file. The metadata fingerprint is
calculated for file size, file ctime, file mtime, file crtime, file
retention time, file uid, file gid, and file type. The fingerprint is
base64 encoded. This field is not included if the scope is ``_data-only_``
.
* Fingerprint Algorithm:

The digest algorithm which is used for the fingerprint computation.
Fingerprint is computed using ``_md5_`` or ``_sha-256_`` digest algorithm.
* Fingerprint Scope:

The scope of the file which is used for the fingerprint computation.
Fingerprint is computed over ``_data-only_`` , ``_metadata-only_`` , or
``_data-and-metadata_`` .
* Fingerprint Start Time:
```

```
The start time of the fingerprint computation in seconds since 1 January
1970 00:00:00 in GMT timezone.
* Formatted Fingerprint Start Time:

The start time of the fingerprint computation in a human-readable format
<day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.
* Fingerprint Version:

The version of the fingerprint output format.
* SnapLock License:

The status of the SnapLock license.
* Vserver UUID:

A universal unique identifier for the Vserver on which the file exists.
* Volume MSID:

The mirror set identifier of the volume where the file resides.
* Volume DSID:

The data set identifier of the volume where the file resides.
* Hostname:

The name of the storage system where the fingerprint operation is
performed.
* Filer ID:

The NVRAM identifier of the storage system.
* Volume Containing Aggregate:

The name of the aggregate in which the volume resides.
* Aggregate ID:

A universal unique identifier for the aggregate containing the volume.
* SnapLock System ComplianceClock:

The System ComplianceClock time in seconds since 1 January 1970 00:00:00
in GMT timezone if it is initialized.
* Formatted SnapLock System ComplianceClock:

The System ComplianceClock time in a human-readable format <day> <month>
<day of month> <hour>:<min>:<sec><year> in GMT timezone if it is
initialized.
* Volume SnapLock Type:

The type of the SnapLock volume. This value is only given for SnapLock
```

```
volumes. Possible values are ``_compliance_`` and ``_enterprise_`` .
* Volume ComplianceClock:

The volume ComplianceClock time in seconds since 1 January 1970 00:00:00
in GMT timezone. This has a value only for SnapLock volumes.
* Formatted Volume ComplianceClock:

The volume ComplianceClock time in a human-readable format <day> <month>
<day of month> <hour>:<min>:<sec><year> in GMT timezone. This has a value
only for SnapLock volumes.
* Volume Expiry Date:

The expiry date of the SnapLock volume in seconds since 1 January 1970
00:00:00 in GMT timezone. The volume expiry date can be in wraparound
format.
* Is Volume Expiry Date Wraparound:

The value is ``_true_`` if the volume expiry date is in wraparound format.
The wraparound format indicates that dates after 19 January 2038 are
mapped from 1 January 1970 through 31 December 2002 to 19 January 2038
through 19 January 2071.
* Formatted Volume Expiry Date:

The expiry date of the SnapLock volume in a human-readable format <day>
<month> <day of month> <hour>:<min>:<sec><year> in GMT timezone. The
volume expiry date can be in wraparound format.
* Filesystem ID:

The filesystem identifier of the volume on which the file resides.
* File ID:

A unique number within the filesystem identifying the file.
* File Type:

The type of the file. Possible values include: ``_worm_`` ,
``_worm_appendable_`` , ``_worm_active_log_`` , ``_worm_log_`` , and
``_regular_`` .
* File Size:

The size of the file in bytes.
* Creation Time:

The creation time of the file in seconds since 1 January 1970 00:00:00 in
GMT timezone.
* Formatted Creation Time:

The creation time of the file in a human-readable format <day> <month>
```

```
<day of month> <hour>:<min>:<sec><year> in GMT timezone.
* Modification Time:

The last modification time of the file in seconds since 1 January 1970
00:00:00 in GMT timezone.
* Formatted Modification Time:

The last modification time of the file in a human-readable format <day>
<month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.
* Changed Time:

The last changed time of the file attributes in seconds since 1 January
1970 00:00:00 in GMT timezone. Time is taken from the system clock for
regular files and from the volume ComplianceClock for WORM files when they
are committed. The changed time can be in wraparound format.
* Is Changed Time Wraparound:

The value is ``_true_`` if the last changed time of the file attributes is
in wraparound format. The wraparound format indicates that dates after 19
January 2038 are mapped from 1 January 1970 through 31 December 2002 to 19
January 2038 through 19 January 2071.
* Formatted Changed Time:

The last changed time of the file attributes in a human-readable format
<day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone. The
changed time can be in wraparound format.
* Retention Time:

The retention time of the files committed to WORM on SnapLock volumes in
seconds since 1 January 1970 00:00:00 in GMT timezone. The retention time
can be in wraparound format.
* Is Retention Time Wraparound:

The value is ``_true_`` if the retention time of the file is in wraparound
format. The wraparound format indicates that dates after 19 January 2038
are mapped from 1 January 1970 through 31 December 2002 to 19 January 2038
through 19 January 2071.
* Formatted Retention Time:

The retention time of the files protected by SnapLock in a human-readable
format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT
timezone. The retention time can be in wraparound format.
* Access Time:

The last access time of the regular files on SnapLock volumes and files on
non-SnapLock volumes attributes in seconds since 1 January 1970 00:00:00
```

```
in GMT timezone.
* Formatted Access Time:

The last access time of the regular files on SnapLock volumes and files on
non-SnapLock volumes attributes in a human-readable format <day> <month>
<day of month> <hour>:<min>:<sec><year> in GMT timezone.
* Owner ID:

The integer identifier of the owner of the file.
* Group ID:

The integer identifier of the group owning the file.
* Owner SID:

The security identifier of the owner of the file when it has NTFS security
style.
* Fingerprint End Time:

The end time of the fingerprint computation in seconds since 1 January
1970 00:00:00 in GMT timezone.
* Formatted Fingerprint End Time:

The end time of the fingerprint computation in a human-readable format
<day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.
* Litigation Count:

The number of litigations on the file.



== 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.

``-session-id <integer>`` - Session ID of Fingerprint Operation::
Specifies the session-id of the fingerprint operation whose output is to
be displayed. It is a unique identifier for the fingerprint operation.
This session-id is returned when the fingerprint operation is started on a
file.
```

```
== Examples

The following example displays the fingerprint information of the
fingerprint session identified by session-id ``_17039367_`` :
```

cluster1::> volume file fingerprint dump -session-id 17039367
Vserver:vs1
Session-ID:17039367
Volume:nfs_slc
Path:/vol/nfs_slc/worm
Data Fingerprint:MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=
Metadata Fingerprint:8iMjqJXiNcqgXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=
Fingerprint Algorithm:SHA256
Fingerprint Scope:data-and-metadata
Fingerprint Start Time:1460612586
Formatted Fingerprint Start Time:Thu Apr 14 05:43:06 GMT 2016
Fingerprint Version:3
SnapLock License:available
Vserver UUID:acf7ae64-00d6-11e6-a027-0050569c55ae
Volume MSID:2152884007
Volume DSID:1028
Hostname:cluster1
Filer ID:5f18eda2-00b0-11e6-914e-6fb45e537b8d
Volume Containing Aggregate:slc_aggr
Aggregate ID:c84634aa-c757-4b98-8f07-eefe32565f67
SnapLock System ComplianceClock:1460610635
Formatted SnapLock System ComplianceClock:Thu Apr 14 05:10:35 GMT 2016
Volume SnapLock Type:compliance
Volume ComplianceClock:1460610635
Formatted Volume ComplianceClock:Thu Apr 14 05:10:35 GMT 2016
Volume Expiry Date:1465880998
Is Volume Expiry Date Wraparound:false
Formatted Volume Expiry Date:Tue Jun 14 05:09:58 GMT 2016
Filesystem ID:1028
File ID:96
File Type:worm
File Size:1048576
Creation Time:1460612515
Formatted Creation Time:Thu Apr 14 05:41:55 GMT 2016
Modification Time:1460612515
Formatted Modification Time:Thu Apr 14 05:41:55 GMT 2016
Changed Time:1460610598
Is Changed Time Wraparound:false
Formatted Changed Time:Thu Apr 14 05:09:58 GMT 2016
Retention Time:1465880998
Is Retention Time Wraparound:false
Formatted Retention Time:Tue Jun 14 05:09:58 GMT 2016
Access Time:-
Formatted Access Time:-
Owner ID:0
Group ID:0
Owner SID:-
Fingerprint End Time:1460612586

```
[[IDfafe9eb7ab4159cd828bedea2e9c37aa]]
= volume file fingerprint show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display fingerprint operation status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file fingerprint show` command returns information for one or
several fingerprint operations. This command requires either `-session-id`
or `-vserver` and `-volume` .

== 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.

``[-session-id <integer>]`` - Session ID of Fingerprint Operation::
If this parameter is specified, the command returns the progress of the
fingerprint operation of the specified session-id. The session-id is a
unique identifier for the fingerprint operation that is returned when the
fingerprint operation is started on a file.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, `-volume` must also be specified. When
queried with `-vserver` and `-volume` , the command returns the progress
of all the fingerprint operations running on that particular volume.

``[-volume <volume name>]`` - Volume Name::
If this parameter is specified, `-vserver` must also be specified. When
```

```
queried with `-vserver` and `-volume` , the command returns the progress
of all the fingerprint operations running on that particular volume.

``[-file </vol/<volume name>/<file path>>]`` - File Path::
If this parameter is specified, the command returns the progress of all
fingerprint operations on the specified file.

``[-operation-status {Unknown|In-Progress|Failed|Aborting|Completed}]`` -
Operation Status::
If this parameter is specified, the command returns the progress of all
fingerprint operations with matching status value.

``[-progress-percentage <integer>]`` - Progress Percentage::
If this parameter is specified, the command returns the progress of all
fingerprint operations with matching progress percentage value.

== Examples

The following example displays the progress of all the fingerprint
operations running on volume ``_nfs_slc_`` :
```

cluster1::> volume file fingerprint show -vserver vs0 -volume nfs_slc
Progress
File-Path Session-ID Status Percentage
--- --- --- ---
/vol/nfs_slc/worm 17104897 Completed 100
/vol/nfs_slc/worm_appedable 17104898 Completed 100
/vol/nfs_slc/regular 17104899 In-Progress 30
3 entries were displayed.

```
[[ID51b29e73c177af0b05df9e7b3262c144]]
= volume file fingerprint start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a file fingerprint computation on a file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file fingerprint start` command starts the fingerprint
computation on a file. The fingerprint computation is started on the file,
```

```
and a session-id is returned. This session-id is an unique identifier for
the fingerprint operation and can be used to get the progress of an
ongoing fingerprint operation as well as the complete fingerprint output
for the file once the operation is completed.

== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the vserver which owns the volume on which the file
resides.

``-file </vol/<volume name>/<file path>>`` - Path::
Specifies the absolute path of the file on which fingerprint needs to be
calculated. The value begins with /vol/<volumename>.

``[-algorithm {MD5|SHA256}]`` - Fingerprint Algorithm::
Specifies the digest algorithm which is used for the fingerprint
computation.
+
Fingerprint can be computed using one of the following digest algorithms:
+
  * ``_md5_``
  * ``_sha-256_``



``[-scope {data-and-metadata|data-only|metadata-only}]`` - Fingerprint
Scope::
Specifies the scope of the file which is used for the fingerprint
computation.
+
Fingerprint can be computed using one of the following scope:
+
  * data-only
  * metadata-only
  * data-and-metadata



== Examples

The following example starts computing fingerprint over data and metadata
for file ``_/vol/nfs_slc/worm_`` using ``_md5_`` hash algorithm. The file
``_/vol/nfs_slc/worm_`` is stored in volume ``_nfs_slc_`` on Vserver
``_vs0_`` .
```

cluster1::> volume file fingerprint start -vserver vs0 -scope data-and-metadata -algorithm md5 -file /vol/nfs_slc/worm
File fingerprint operation is queued. Run "volume file fingerprint show -session-id 16973825" to view the fingerprint session status.

```
[[ID5be8c5ff7ee153f35749137263102a3b]]
= volume file retention show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display retention time of a file protected by SnapLock.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume file retention show` command displays the retention time of a
file protected by SnapLock given `-vserver` and `-file` .

== 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.

``[-vserver <vserver name>]`` - Vserver Name::
This specifies the Vserver on which the volume is located having the file
in a SnapLock volume. An exact value is required for this parameter.

``[-file </vol/<volume name>/<file path>>]`` - File path::
This specifies the absolute path of the file in a SnapLock volume. The
value begins with /vol/<volumename>. An exact value is required for this
parameter.

``[-retention-time <integer>]`` - Retention Time of the File::
If this parameter is specified, the command displays information only
about the file protected by SnapLock that matches the specified
`retention-time` . The value represents file retention time in seconds
```

```
since 1 January 1970 00:00:00.

``[-formatted-retention-time <text>]`` - Formatted Retention Period::
If this parameter is specified, the command displays information only
about the file protected by SnapLock that matches the specified
`formatted-retention-time` . The expiry date format is <day> <month> <day
of month> <hour>:<min>:<sec><year> in GMT timezone taking care of
wraparound. A value of ``_infinite_`` indicates that this file has
infinite retention time. A value of ``_indefinite_`` indicates that this
file has indefinite retention time. A value of ``_unspecified_`` indicates
that this file will be retained forever; however, the retention time can
be changed to an absolute value.

``[-is-wraparound {true|false}]`` - Is Retention Time Wraparound::
If this parameter is specified, the command displays information only
about the file protected by SnapLock that matches the specified `-is-
wraparound` . The value is `true` if the date represented in retention
time is in wraparound format. The wraparound format indicates that dates
after 19 January 2038 are mapped from 1 January 1970 through 31 December
2002 to 19 January 2038 through 19 January 2071.

``[-seconds-until-expiry <integer>]`` - Seconds until the File Expires::
If this parameter is specified, the command displays information only
about the file protected by SnapLock that matches the specified `-seconds-
until-expiry` . The value represents the number of seconds until the
expiration time.

``[-is-expired {true|false}]`` - Has the File Expired::
If this parameter is specified, the command displays information only
about the file protected by SnapLock that matches the specified `-is-
expired` . The value represents file expiration status. The value is
``_false_`` if the file is under active retention or ``_true_`` if the
file is past its expiry time.

== Examples

The following example displays the retention time of the file
``_/vol/nfs1/file1_`` :
```

cluster1::> volume file retention show -vserver vs0 -file /vol/nfs1/file1
Vserver : vs0
Path : /vol/nfs1/file1
Retention Time (Secs from Epoch) : 1439111404
Formatted Retention Time : Sun Aug 9 09:10:04 GMT 2015
Is Retention Time Wraparound : false

```
:leveloffset: -1


= volume flexcache commands

:leveloffset: +1


[[ID8e9e1ee9d44c5aab7302b3d683e67555]]
= volume flexcache config-refresh
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Refresh FlexCache configuration for a peer volume

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

 The `volume flexcache config-refresh` command is used to refresh the
FlexCache configuration. It can be used to update the configuration on
either the FlexCache or origin of a FlexCache cluster. This command only
needs to be run if the automatic refresh failed and the corresponding EMS
was generated.

NOTE: This command must be run on the peer cluster. For example, refresh
of origin of a FlexCache volume must be run from the FlexCache cluster.
Refresh a FlexCache volume must be run from the of an origin of a
FlexCache cluster.



== Parameters

``-peer-vserver <vserver name>`` - Peer Vserver [.advaced-
priv]#(privilege: advanced)#::
Name of the Vserver for which the configuration is being refreshed.

``-peer-volume <volume name>`` - Peer Volume [.advaced-priv]#(privilege:
advanced)#::
Name of the volume for which the configuration is being refreshed.

``-peer-endpoint-type {cache|origin}`` - Origin/Cache Volume [.advaced-
```

```
priv]#(privilege: advanced)#::
The peer-endpoint-type specifies the FlexCache endpoint type of the peer
volume. Possible values are ``_cache_`` for FlexCache volumes and
``_origin_`` for origin of a FlexCache volumes.

== Examples

The following example triggers config-refresh on origin of a FlexCache
volume "origin1".
```

cluster1::> flexcache config-refresh -peer-vserver vs34 -peer-volume origin1 -peer-endpoint-type origin
(volume flexcache config-refresh)

```
The following example triggers config-refresh on FlexCache volume "fc1".
```

cluster1::> flexcache config-refresh -peer-vserver vs34 -peer-volume fc1 -peer-endpoint-type cache
(volume flexcache config-refresh)

```
The following example triggers config-refresh on FlexCache volume "fc1"
with an incorrect peer-endpoint-type.
```

cluster1::> flexcache config-refresh -peer-vserver vs34 -peer-volume fc1 -peer-endpoint-type origin
(volume flexcache config-refresh)
Error: command failed: Failed to store the configuration for peer volume "fc1"
in Vserver "vs34" on cluster "cluster1". Check the FlexCache
configuration on the local cluster and retry the operation.

```
[[ID66e0d2f354389cfa98058e14496827e1]]
= volume flexcache create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new cache relationship

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 The `volume flexcache create` command is used to create a FlexCache
volume. It also creates the relationship between the FlexCache volume and
```

```
the specified origin volume.

NOTE:

* If the `vserver` and `origin-vserver` are different, the Vservers must
be in a peer relationship.
*





== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This specifies the Vserver on which the FlexCache volume is to be created.

``-volume <volume name>`` - Cache Volume Name::
This specifies the name of the FlexCache volume that is to be created.

{ ``[-aggr-list <aggregate name>,...]`` - List of Aggregates for FlexGroup
Constituents::
Specifies an array of names of aggregates to be used for creating the
FlexCache volume. Each entry in the list will create a constituent on the
specified aggregate. The root constituent will always be placed on the
first aggregate in the list, unless `optimize-aggr-list` is specified as
``_true_`` . An aggregate may be specified multiple times to have multiple
constituents created on it. This parameter only applies to FlexGroup
volumes.

``[-aggr-list-multiplier <integer>]`` - Aggregate List Repeat Count::
Specifies the number of times to iterate over the aggregates listed with
the `-aggr-list` parameter when creating a FlexGroup. The aggregate list
will be repeated the specified number of times. Example:
```

-aggr-list aggr1,aggr2 -aggr-list-multiplier 2

```
will cause four constituents to be created in the order ``_aggr1_`` ,
``_aggr2_`` , ``_aggr1_`` , ``_aggr2_`` .
+
The default value is 4.

``[-optimize-aggr-list {true|false}]`` - Have the System Optimize the
Order of the Aggregate List::
Specifies whether to create the constituents of the FlexCache volume on
the aggregates specified in the `aggr-list` in the order they are
```

```
specified, or whether the system should optimize the ordering of the
aggregates. If this value is ``_true_`` , the system will optimize the
ordering of the aggregates specified in the `aggr-list` . If this value is
``_false_`` the order of the `aggr-list` will be unchanged. The default
value is ``_false_`` . This parameter only applies to FlexGroup volumes.

| ``[-auto-provision-as <FlexGroup>]`` - Automatically Provision as Volume
of Type::
Use this parameter to automatically select existing aggregates for
provisioning the FlexCache volume.

``[-use-tiered-aggregate {true|false}]`` - Automatically Provision
FlexGroup on FabricPool }::
This parameter specifies whether or not FabricPool-enabled aggregates are
selected when provisioning a FlexCache volume using the `auto-provision-
as` parameter. Only FabricPool-enabled aggregates are used if this
parameter is set to true and only non-FabricPool-enabled aggregates are
used if this parameter is set to false. The default value is ``_false_`` .
Tiering-policy is always ``_none_`` for FlexCache.
+
This parameter only applies to flexible caches created using the `-auto-
provision-as` parameter.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Size::
This optionally specifies the size of the FlexCache volume. The size is
specified as a number followed by a unit designation: k (kilobytes), m
(megabytes), g (gigabytes), or t (terabytes). If the unit designation is
not specified, bytes are used as the unit, and the specified number is
rounded up to the nearest 4 KB. If the size parameter in not specified, it
defaults to 10% of the origin volume size.

``[-origin-vserver <vserver name>]`` - Origin Vserver Name::
 This specifies the name of the Vserver where the origin volume is
located.

``-origin-volume <volume name>`` - Origin Volume Name::
 This specifies the name of the origin volume.

``[-junction-path <junction path>]`` - Cache Junction Path::
This optionally specifies the FlexCache volume's junction path. The
junction path name is case insensitive and must be unique within a
Vserver's namespace.

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). When set to
```

```
true, the command will not return until the operation completes.

``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
This optionally specifies the space guarantee style for the volume. A
value of ``_volume_`` reserves space on the aggregates for the entire
volume. A value of ``_none_`` reserves no space on the aggregates, meaning
that writes can fail if an aggregate runs out of space. The default
setting for the FlexCache volumes is ``_none_`` .

``[-is-cifs-enabled {true|false}]`` - CIFS access Enabled/Disabled
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies CIFS access for the volume. The default value is
``_true_`` .

``[-preserve-msid {true|false}]`` - Preserve Origin MSID on the Cache
[.advaced-priv]#(privilege: advanced)#::
Create the cache with the same MSID as the origin. Set this parameter to
``_true_`` to create FlexCache DR. The default setting for the FlexCache
volumes is ``_false_`` to avoid a situation where a FlexCache create
operation using `-preserve-msid` fails on a Vserver because of collision
with a previously used MSID.

``[-is-global-file-locking-enabled {true|false}]`` - Global File Locking
Mode Enabled/Disabled::
This parameter specifies the global file locking mode for the volume. The
default value of global file locking mode is ``_false_`` .

``[-is-nfsv4-enabled {true|false}]`` - NFSv4 access Enabled/Disabled
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies NFSv4 access for the volume. The default value is
``_true_`` .

``[-is-writeback-enabled {true|false}]`` - Is Writeback Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether writeback is enabled for the FlexCache
volume. Writeback is a storage method where data is first written to the
FlexCache volume and then written to the origin of a FlexCache volume. The
default value is ``_false_`` .

``[-is-relative-size-enabled {true|false}]`` - Is Relative Sizing
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether the relative sizing is enabled for the
FlexCache volume. Relative sizing is introduced as a part of follow the
origin feature. When relative sizing is enabled, it will block any
modifications done manually in the absolute size of the FlexCache. The
size of FlexCache will be calculated and entered automatically based on
```

```
the size of the origin. The default value for create is ``_false_`` .

``[-relative-size-percentage <integer>]`` - Relative FlexCache Size
Percentage [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the percent size FlexCache should have relative
to the total size of the origin. The default value for create is 10%.

``[-override-encryption {true|false}]`` - Override FlexCache Volume
Encryption::
Use this parameter to create a plaintext FlexCache volume of an encrypted
origin volume. The default value is ``_false_`` .
+
NOTE: This parameter cannot be modified after the FlexCache volume is
created.
+



``[-is-atime-scrub-enabled {true|false}]`` - Is atime-Based Scrubbing
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether automatic scrubbing of inactive files
based on atime is enabled. When enabled, inactive files are automatically
scrubbed based on specified atime duration. The scrubbing operation will
only occur if `-atime-update` is enabled on the FlexCache volume. The
default value for the atime scrubber is ``_false_`` .

``[-atime-scrub-period <integer>]`` - Days After Which An Inactive File Is
Scrubbed [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the number of days for inactivity on a file after
which the file will be scrubbed out. The default value is 30 days.

``[-is-cifs-change-notify-enabled {true|false}]`` - Is CIFS CHANGE_NOTIFY
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether the FlexCache volume sends change
notifications to CIFS clients. The default value is ``_false_`` .

== Examples

The following example triggers FlexCache volume create:
```

cluster1::> flexcache create -vserver vs34 -volume fc1 -aggr-list aggr34,aggr43 -origin-volume origin1 -size 400m
(volume flexcache create)
[Job 894] Job succeeded: Successful

```
cluster1::> flexcache create -vserver vs34 -volume fc3 -auto-provision-as flexgroup -origin-volume origin1 -size
400m
(volume flexcache create)
[Job 898] Job succeeded: Successful
```

```
cluster1::> flexcache create -vserver vs34 -volume fc4 -aggr-list aggr34,aggr43 -origin-volume origin1 -size
400m -junction-path /fc4
(volume flexcache create)
[Job 903] Job succeeded: Successful
```

```
[[ID977b3a0378dc9d6a355a8aebd2a5679b]]
= volume flexcache delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a cache relationship

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume flexcache delete` command deletes the specified FlexCache
volumes and their relationships.

NOTE: FlexCache volumes must be offline (see xref:{relative_path}volume-
offline.html[volume offline] ) to be deleted.



== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This specifies the name of the Vserver from which the FlexCache volume is
to be deleted.

``-volume <volume name>`` - Cache Volume Name::
 This specifies the name of the FlexCache volume that is to be deleted.

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the operation runs in the foreground. The default
setting is ``_true_`` (the operation runs in the foreground). When set to
true, the command will not return until the operation completes.

== Examples

The following example deletes FlexCache volume "fc1":
```

cluster1::> flexcache delete -volume fc1 -vserver vs34
(volume flexcache delete)

Error: command failed: Volume fc1 in Vserver vs34 must be offline to be
deleted. Use "volume offline -vserver vs34 -volume fc1" command to

offline the volume

```
cluster1::> volume offline -vserver vs34 -volume fc1
Volume "vs34:fc1" is now offline.
```

```
cluster1::> flexcache delete -volume fc1 -vserver vs34
(volume flexcache delete)
[Job 891] Job succeeded: Successful
```

```
cluster1::> flexcache delete -volume fc1 -vserver vs34 -foreground false
(volume flexcache delete)
```

```
== Related Links

* xref:{relative_path}volume-offline.html[volume offline]




[[IDff2e5a6f22221433d4004674d93024e1]]
= volume flexcache prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disables FlexCache features not supported by a previous version

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command disables the FlexCache feature enabled for a version.

== Parameters

``[-disable-feature-set <ClusterVersion>]`` - Cluster Version [.advaced-
priv]#(privilege: advanced)#::
 This specifies the version for which we are running this command.

== Examples

The following example disables FlexCache features for cluster "cluster1":
```

cluster1::> flexcache prepare-to-downgrade -disable-feature-set 9.8.0
FlexCache features introduced in 9.8 are successfully disabled.

```
[[IDc6a7406d5c4193b05a7497b736f93eff]]
= volume flexcache show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Display cache relationships

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume flexcache show` command displays information about FlexCache
volumes. The command output depends on the parameter or parameters
specified with the command. If no parameters are specified, the command
displays the following information about all FlexCache volumes:

* Vserver name
* Volume name
* Size
* Space guarantee
* Origin Vserver
* Origin volume
* Origin cluster



To display detailed information about all FlexCache volumes, run the
command with the `-instance` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
If this parameter is specified, the command displays information about all
values.

``[-vserver <Vserver Name>]`` - Vserver Name::
If this parameter and the `-volume` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
FlexCache volumes on the specified Vserver.

``[-volume <volume name>]`` - Cache Volume Name::
If this parameter is specified, the command displays detailed information
about the specified FlexCache volume.

``[-aggr-list <aggregate name>,...]`` - List of Aggregates for FlexGroup
Constituents::
If this parameter is specified, the command displays information only
```

```
about the FlexCache volume or volumes that are located on the specified
list of storage aggregates.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Size::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that have the specified size. Size
is the maximum amount of space a volume can consume from its associated
aggregate(s), including user data, metadata, snapshots, and snapshot
reserve. Note that for volumes without a `-space-guarantee` of
``_volume_`` , the ability to fill the volume to this maximum size depends
on the space available in the associated aggregate or aggregates.

``[-flexgroup-msid <integer>]`` - Cache FlexGroup MSID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that have the specified FlexGroup
master data-set ID.

``[-origin-vserver <vserver name>]`` - Origin Vserver Name::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes which have a relationship with the
specified origin-vserver.

``[-origin-vserver-uuid <UUID>]`` - Origin Vserver UUID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes which have a relationship with the
origin-vserver UUID.

``[-origin-volume <volume name>]`` - Origin Volume Name::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that have a relationship with the
specified origin-volume.

``[-origin-volume-msid <integer>]`` - Origin Volume MSID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that have a relationship with the
specified origin master data set ID.

``[-origin-cluster <Cluster name>]`` - Origin Cluster Name::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that have a relationship with the
specified origin-cluster.

``[-junction-path <junction path>]`` - Cache Junction Path::
```

```
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that have the specified junction
path.

``[-create-time <Date>]`` - FlexCache Create Time::
If you specify this parameter, the command displays information only about
the FlexCache volume or volumes for which the `create-time` option matches
the specified input.

``[-s, -space-guarantee {none|volume}]`` - Space Guarantee Style::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that have the specified space
guarantee style.

``[-relationship-type {Intravserver|Intervserver|Intercluster}]`` -
Relationship Type [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the relationship type.

``[-is-cifs-enabled {true|false}]`` - CIFS access Enabled/Disabled
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the value for CIFS access.

``[-preserve-msid {true|false}]`` - Preserve Origin MSID on the Cache
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that match the given value for the
`-preserve-msid` . If this parameter is ``_true_`` , then FlexCache DR is
enabled for the FlexCache volume.

``[-is-global-file-locking-enabled {true|false}]`` - Global File Locking
Mode Enabled/Disabled::
If this parameter is specified, the command displays the FlexCache volumes
matching the global file locking mode.

``[-is-nfsv4-enabled {true|false}]`` - NFSv4 access Enabled/Disabled
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the value for NFSv4 access.

``[-is-writeback-enabled {true|false}]`` - Is Writeback Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether writeback is enabled for the FlexCache
volume. Writeback is a storage method where data is first written to the
FlexCache volume and then written to the origin of a FlexCache volume.
```

```
``[-is-relative-size-enabled {true|false}]`` - Is Relative Sizing
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that match the given value for `-is-
relative-size-enabled` . When relative size is enabled, the FlexCache size
will be automatically updated based on the size of the origin.

``[-relative-size-percentage <integer>]`` - Relative FlexCache Size
Percentage [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the information only
about the FlexCache volume or volumes for which the value of `-relative-
size-percentage` is specified. The relative size percentage is used to
calculate the size of a FlexCache volume when the is-relative-size-enabled
field is ``_true_`` .

``[-override-encryption {true|false}]`` - Override FlexCache Volume
Encryption::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that match the given value for the
`-override-encryption` . If this parameter is ``_true_`` , then the volume
is a plaintext FlexCache volume.

``[-is-atime-scrub-enabled {true|false}]`` - Is atime-Based Scrubbing
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that match the given value for `-is-
atime-scrub-enabled` . When atime scrubber is enabled, the FlexCache
volume will automatically scrub inactive files based on the specified
atime duration.

``[-atime-scrub-period <integer>]`` - Days After Which An Inactive File Is
Scrubbed [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes for which the value of `-atime-
scrub-period` is specified. The atime scrub period is used to remove the
files from the FlexCache volume when the duration is met. The default
duration to scrub files is set to ``_30 days_`` .

``[-is-cifs-change-notify-enabled {true|false}]`` - Is CIFS CHANGE_NOTIFY
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the FlexCache volume or volumes that match the given value for `-is-
cifs-change-notify-enabled` . When CIFS CHANGE_NOTIFY is enabled, the
FlexCache volume will notify clients about changes on a directory.

== Examples
```

```
The following example displays information about all FlexCache volumes on
the Vserver named "vs34":
```

cluster1::> flexcache show -vserver vs34
(volume flexcache show)
Vserver Volume Size Origin-Vserver Origin-Volume Origin-Cluster
------- ----------- ---------- -------------- ------------- --------------
vs34 fc1 800MB vs34 origin1 cluster-2
vs34 fc2 800MB vs34 origin1 cluster-2
2 entries were displayed.

```
The following example displays detailed information about a FlexCache
volume named fc1 on an SVM named vs34:
```

cluster1::> flexcache show -vserver vs34 -volume fc1 -instance
(volume flexcache show)
Vserver: vs34
Cache Volume Name: fc1
List of Aggregates for FlexGroup Constituents: aggr34
Volume Size: 800MB
Cache Flexgroup MSID: 2155934574
Origin Vserver Name: vs34
Origin Vserver UUID: a8717aeb-2826-11e8-bf56-00505695f37a
Origin Volume Name: origin1
Origin Volume MSID: 2155934545
Origin Cluster Name: cluster-2
Cache Junction Path: -
FlexCache Create Time: Thu Aug 23 04:36:19 2018
Relationship Type: inter-vserver
Preserve Origin MSID on the Cache: false

```
[[IDaa9062126496c637f4d4378a56b48aec]]
= volume flexcache sync-properties
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Sync volume properties

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume flexcache sync-properties` command is used to sync a FlexCache
volume's properties with its origin volume or to update FlexCache origin
```

```
volume name in FlexCache relationship during a origin volume rename
operation in disconnected mode.

NOTE:

* This command should be used when there is a failure to sync properties
from the origin of a FlexCache volume to the FlexCache volume or to update
the FlexCache origin volume name in the FlexCache relationship in the case
of a volume rename from the origin in disconnected mode.
*





== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the SVM on which the volume is located.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
Specifies the volume name whose properties need to be synced. If the
volume is a FlexCache volume, its properties are synced with its origin
volume. If the volume is a FlexCache origin volume, the origin volume name
is updated in the FlexCache relationship where the origin volume is
renamed in disconnected mode.

== Examples
```

cluster1::> flexcache sync-properties -volume fc1 -vserver vs34
(volume flexcache sync-properties)

```
[[ID11ac4e31bd94eacd313e73f688a70a97]]
= volume flexcache config modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify cache configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.
```

```
== Description

 The `volume flexcache config modify` command is used to modify the
configuration for a FlexCache volume.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver in which the FlexCache volume exists.

``-volume <volume name>`` - Cache Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the FlexCache volume.

``[-is-cifs-enabled {true|false}]`` - Is CIFS Access Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether CIFS access is allowed on a FlexCache
volume.

``[-is-nfsv4-enabled {true|false}]`` - Is NFSv4 Access Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether NFSv4 access is allowed on a FlexCache
volume.

``[-is-writeback-enabled {true|false}]`` - Is Writeback Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether writeback is enabled for the FlexCache
volume. Writeback is a storage method where data is first written to the
FlexCache volume and then written to the origin of a FlexCache volume.

``[-is-relative-size-enabled {true|false}]`` - Is Relative Sizing Enabled
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether the relative sizing is enabled for the
FlexCache volume. Relative sizing is introduced as a part of follow the
origin feature. When relative sizing is enabled, it will block any
modifications done manually in the absolute size of the FlexCache. The
size of FlexCache will be calculated and entered automatically based on
the size of the origin. The default value for create is ``_false_`` .

``[-relative-size-percentage <integer>]`` - Relative FlexCache Size
Percentage [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the percent size FlexCache should have relative
to the total size of the origin. The default value for create is 10%.

``[-is-atime-scrub-enabled {true|false}]`` - Is atime based Scrubbing
```

```
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies if the files can be scrubbed out automatically if
they have not been accessed for a specified duration. This feature to
scrub out inactive files is dependent on enabling the atime feature on the
FlexCache volume. The default value for create is ``_disabled_`` .

``[-atime-scrub-period <integer>]`` - Days after which an inactive file is
scrubbed [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the duration in days after which the inactive
files can be scrubbed out from the cache volume. The default value for
create is ``_30 days_`` when atime-scrub is enabled.

``[-is-cifs-change-notify-enabled {true|false}]`` - Is CIFS CHANGE_NOTIFY
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether the FlexCache volume sends change
notifications to CIFS clients. The default value is ``_false_`` .

== Examples

The following example triggers FlexCache volume configuration modify:
```

cluster1::> volume flexcache config modify -vserver vs34 -volume fc1 -is-cifs-enabled true

```

```

cluster1::> volume flexcache config modify -volume originvol -is-cifs-enabled false
(volume flexcache config modify)
Error: command failed: FlexCache config modify is only applicable for volumes
with "flexcache-endpoint-type" cache.

```
[[ID9ac4b44ddbe399cc0e1bb776de83c31c]]
= volume flexcache config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display cache configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

 The `volume flexcache config show` command is used to view the FlexCache
```

```
configuration for a FlexCache volume.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-volume` parameter are specified, the command
displays config information about the specified volume. If this parameter
is specified by itself, the command displays information about FlexCache
volumes on the specified Vserver.

``[-volume <volume name>]`` - Cache Volume Name [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays config information
about the specified FlexCache volumes matching the specified name.

``[-is-cifs-enabled {true|false}]`` - Is CIFS Access Enabled [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the value for CIFS access.

``[-is-nfsv4-enabled {true|false}]`` - Is NFSv4 Access Enabled [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the value for NFSv4 access.

``[-is-writeback-enabled {true|false}]`` - Is Writeback Enabled [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the writeback mode. Writeback is a storage method where data is
first written to the FlexCache volume and then written to the origin of a
FlexCache volume.

``[-is-relative-size-enabled {true|false}]`` - Is Relative Sizing Enabled
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the value for relative size.
```

```
``[-relative-size-percentage <integer>]`` - Relative FlexCache Size
Percentage [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the percentage for relative size.

``[-is-atime-scrub-enabled {true|false}]`` - Is atime based Scrubbing
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the atime scrubbing mode.

``[-atime-scrub-period <integer>]`` - Days after which an inactive file is
scrubbed [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the period for atime scrubbing.

``[-is-cifs-change-notify-enabled {true|false}]`` - Is CIFS CHANGE_NOTIFY
Enabled/Disabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the FlexCache volumes
matching the value for CIFS CHANGE_NOTIFY mode.

== Examples

The following example triggers FlexCache volume configuration show:
```

cluster1::> flexcache config show
(volume flexcache config show)
Vserver Volume CIFS Access Enabled
------- ---------- ------------------
vs1 cachevol false
vs1 cachevol_2 true
2 entries were displayed.

```
[[ID2397975702e433c7c8d0b0bb08368e28]]
= volume flexcache origin show-caches
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display all the caches connected to origin

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `volume flexcache origin show-caches` command displays FlexCache
relationships for origin of a FlexCache volumes on the origin cluster.

== 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.

``[-origin-vserver <Vserver Name>]`` - Vserver Name::
If this parameter and the `origin-volume` parameter are specified, the
command displays the FlexCache relationship for the specified origin of a
FlexCache volume. If this parameter is specified by itself, the command
displays all the FlexCache relationships for all origin of a FlexCache
volumes in the specified Vserver.

``[-origin-volume <volume name>]`` - Origin Volume::
If this parameter is specified, the command displays all the FlexCache
relationships for the specified origin of a FlexCache volume.

``[-cache-vserver <vserver name>]`` - Cache Vserver::
If this parameter is specified then the command displays FlexCache
relationships for the specified Vserver hosting FlexCache volumes.

``[-cache-volume <volume name>]`` - Cache Volume::
If this parameter is specified then the command displays FlexCache
relationships for the specified FlexCache volume.

``[-cache-cluster <Cluster name>]`` - Cache Cluster::
If this parameter is specified then the command displays FlexCache
relationships for the specified cluster hosting FlexCache volumes.

``[-cache-volume-msid <integer>]`` - Cache Volume MSID::
If this parameter is specified then the command displays FlexCache
relationships for the specified FlexCache volume's FlexGroup master data-
set ID.

``[-relationship-create-time <Date>]`` - Relationship Create Time::
If you specify this parameter, the command displays FlexCache
relationships for which the `relationship-create-time` option matches the
specified input.
```

```
``[-relationship-type {Intravserver|Intervserver|Intercluster}]`` -
Relationship Type [.advaced-priv]#(privilege: advanced)#::
 If this parameter is specified, the command displays the origin of a
FlexCache volume matching the relationship type

``[-preserve-msid {true|false}]`` - Preserve Origin MSID on the Cache
[.advaced-priv]#(privilege: advanced)#::
 If this parameter is specified, the command displays the origin of a
FlexCache volume whose Msid is preserved

== Examples

 The following example displays information about all origin of a
FlexCache volumes on the Vserver named vs34:
```

cluster1::> flexcache origin show-caches -origin-vserver vs34
(volume flexcache origin show-caches)
Origin-Vserver Origin-Volume Cache-Vserver Cache-Volume Cache-Cluster
--- --- --- --- ---
vs34 origin1 vs56 fc1_c3_origin1
cluster-3
vs34 origin1 vs34 fc1 cluster-2
vs34 origin1 vs34 fc2 cluster-2
vs34 origin2_new vs56 fc1_c3_origin2
cluster-3
4 entries were displayed.

```

```

cluster1::> flexcache origin show-caches -origin-vserver vs56 -instance
(volume flexcache origin show-caches)
Vserver: vs56
Origin Volume: origin
Cache Vserver: vs56
Cache Volume: fc1
Cache Cluster: cluster-3
Cache Volume MSID: 2156002002
Relationship Create Time: Thu Aug 23 04:36:24 2018
Relationship Type: intra-vserver

```
[[ID3eeebdc4b72954f83a5660425483fd97]]
= volume flexcache origin config modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify FlexCache Origin Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume flexcache origin config modify` command modifies volume
specific options of an origin of a FlexCache volume.

== Parameters

``-origin-vserver <Vserver Name>`` - Vserver Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the Vserver in which the origin of a FlexCache volume is
located.

``-origin-volume <volume name>`` - Origin Volume Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the origin of a FlexCache volume.

``[-is-bli-enabled {true|false}]`` - Is BLI Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether data invalidation at block level is
enabled or not. By default, `-is-bli-enabled` is ``_false_`` .

``[-is-global-file-locking-enabled {true|false}]`` - Is Global File
Locking Enabled [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the global file locking mode for the volume. By
default, `-is-global-file-locking-enabled` is ``_false_`` . When global
file locking mode is enabled for the volume, the
"is_disconnected_mode_off_for_locks" flag is always set to ``_true_`` .

== Examples

The following example modifies the is-bli-enabled field value to true:
```

cluster1::> flexcache origin config modify -origin-volume origin_flexgroup -origin-vserver vs3_c2 -is-bli-enabled true
(volume flexcache origin config show)

```
```

cluster1::> flexcache origin config show -origin-volume origin_flexgroup -origin-vserver vs3_c2
(volume flexcache origin config show)
Vserver: vs3_c2
Origin Volume Name: origin_flexgroup
Is BLI Enabled: true

```
[[IDad51bf68e4dbefeb7f1e71d91905364e]]
= volume flexcache origin config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show FlexCache Origin Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume flexcache origin config show` command displays volume specific
options of an origin of a FlexCache volume.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-origin-vserver <Vserver Name>]`` - Vserver Name [.advaced-
priv]#(privilege: advanced)#::
If this parameter and the `-origin-volume` parameter are specified, the
command displays options for the specified origin of a FlexCache volume.
If this parameter is specified by itself, the command displays options for
all origin of a FlexCache volumes in the specified Vserver.
```

```
``[-origin-volume <volume name>]`` - Origin Volume Name [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays options for the
specified origin of a FlexCache volume.

``[-is-bli-enabled {true|false}]`` - Is BLI Enabled [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the origin of a
FlexCache volume matching the specified value. Possible values are:
+
  * ``_false_`` - Block level invalidation is not set
  * ``_true_`` - Block level invalidation is set



``[-is-global-file-locking-enabled {true|false}]`` - Is Global File
Locking Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays the origin
configurations matching the specified value. Global file locking mode is a
mode where protocol read locking semantics are enforced across all
FlexCache volumes and origins of a FlexCache volume. To completely
disable/enable global file locking mode, use `flexcache origin config
modify` in advanced privilege with `-is-global-file-locking-enabled` set
to ``_true_`` /``_false_`` depending upon the requirement. When global
file locking mode is enabled for the volume, the
"is_disconnected_mode_off_for_locks" flag is always set to ``_true_`` .
Possible values are:
+
  * ``_false_`` - Global file locking mode is not set
  * ``_true_`` - Global file locking mode is set



== Examples

The following example displays FlexCache origin options:
```

```
cluster1::> volume flexcache origin config show
(volume flexcache origin config show)
Origin-Vserver Origin-Volume Is-BLI-Enabled Is-Global-File-Locking-
Enabled
--------------- --------------- -------------- ------------------------------
vs34 origin1 true true
vs34 origin2_new false false
2 entries were displayed.
```

```
[[ID557e88cb55d91ed4626d4f049592077d]]
= volume flexcache prepopulate start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Prepopulate flexcache volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

 The `volume flexcache prepopulate start` command is used to prepopulate a
FlexCache volume with the data of given directory path(s) from the origin
volume.

NOTE:  The FlexCache prepopulate job fails when the directories listed
with the `-path-list` have one or more non-existing paths.



== Parameters

``-cache-vserver <vserver name>`` - Cache Vserver Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the Vserver that has the FlexCache volume to be
prepopulated.

``-cache-volume <volume name>`` - Cache Volume Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the FlexCache volume which needs to be prepopulated.

``-path-list <text>,...`` - List of Paths [.advaced-priv]#(privilege:
advanced)#::
This specifies an array of directory names that need to be prepopulated in
the FlexCache volume.

``[-isRecursion {true|false}]`` - Recursively populate [.advaced-
priv]#(privilege: advanced)#::
 This specifies that the directories listed with the `-path-list` need to
be recursively prepopulated in the `-cache-volume` FlexCache volume.
+
The default value is true.
```

```
``[-exclude-path-list <text>,...]`` - Exclude List of Paths [.advaced-
priv]#(privilege: advanced)#::
This specifies an array of directory names that need not be populated.
+
NOTE:  The path provided should be absolute path. If same path provided is
for include and exclude path list then include-path will take the
preference
+


== Examples

The following example triggers FlexCache volume prepopulate:
```

cluster1::> flexcache prepopulate start -cache-volume vol_cache2_vs3_c2_vol_origin1_vs1_c1 -cache-vserver vs3_c2 -path-list /dir1 -isRecursion false
(volume flexcache prepopulate start)
[JobId 164]: FlexCache prepopulate job queued.

```
```

cluster1::> flexcache prepopulate start -cache-volume vol_cache2_vs3_c2_vol_origin1_vs1_c1 -cache-vserver vs3_c2 -path-list /dir1
(volume flexcache prepopulate start)
[JobId 188]: FlexCache prepopulate job queued.

```
```

cluster1::> flexcache prepopulate start -cache-volume vol_cache2_vs3_c2_vol_origin1_vs1_c1 -cache-vserver vs3_c2 -path-list /dir1, dir5, dir6
(volume flexcache prepopulate start)

Error: command failed: Path(s) "dir5, dir6" does not exist in origin volume
"vol_origin1_vs1_c1" in Vserver "vs1_c1".

```
:leveloffset: -1


= volume flexgroup commands

:leveloffset: +1


[[IDcbdaadac9a216101af3229bc8c50ddee]]
= volume flexgroup qtree-disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable qtree support on a FlexGroup

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume flexgroup qtree-disable` command disables qtree support on a
FlexGroup.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver in which the FlexGroup is located.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the FlexGroup.

== Examples

The following example disables the qtree support on a FlexVol named "fg"
in Vserver "vs0":
```

cluster::*> volume flexgroup qtree-disable -vserver vs0 -volume fg

```
:leveloffset: -1


= volume inode-upgrade commands

:leveloffset: +1


[[ID73a4101c16dd20319ca491d0eb2faf53]]
= volume inode-upgrade resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume suspended inode upgrade

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume inode-upgrade resume` command resumes suspended inode upgrade
process. The inode upgrade process may be suspended earlier due to
performance reasons.

== Parameters

``-vserver <vserver name>`` - VServer Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume for which the inode upgrade process is to be
resumed.

== Examples

The following example resumes a volume upgrade process.
```

cluster1::> volume inode-upgrade resume -vserver vs0 -volume vol1

```
[[ID85dd22f0e47802229fe0c85d4a647a5b]]
```

```
= volume inode-upgrade show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Inode Upgrade Progress

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume inode-upgrade show` command displays information about volumes
in the middle of the inode upgrade process. The command output depends on
the parameter or parameters specified with the command. If no parameters
are specified, the command displays the default fields about all volumes
in the middle of the inode upgrade process. Default fields are `vserver` ,
`volume` , `aggregate` , `status` , `scan-percent` , `remaining-time` ,
`space-needed` , and `scanner-progress` .

== Parameters

``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is specified, the command displays information about all
entries.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-volume` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-vserver` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
all volumes that match the specified name.

``[-node <nodename>]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
```

```
If this parameter is specified, the command displays information only
about the volume or volumes that are located on the specified storage
system.

``[-vol-dsid <integer>]`` - Volume DSID [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified data set ID.

``[-vol-uuid <UUID>]`` - Volume UUID [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified UUID.

``[-volume-msid <integer>]`` - Volume MSID [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified master data set ID.

``[-vserver-uuid <UUID>]`` - Vserver UUID [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the volume on the Vserver that has the specified UUID.

``[-aggregate <aggregate name>]`` - Aggregate Name [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that are located on the specified storage
aggregate.

``[-aggregate-uuid <UUID>]`` - Aggregate UUID [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that are located on the storage aggregate with
the specified UUID.

``[-status {pending|scanning|suspended-initalizing|suspended|cleanup-
pending|cleanup|cleanup-done|suspended-aborting|suspended-
removing|suspended-while-removing|suspended-ironing}]`` - Upgrade Status
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified inode upgrade status.

``[-scan-percent <percent>]`` - Upgrade Scan Percent Complete [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
```

```
about the volume or volumes that match the specified inode upgrade
progress percentage.

``[-space-needed {<integer>[KB|MB|GB|TB|PB]}]`` - Space Needed to Complete
Upgrade [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes where the space needed to complete the upgrade
process match the specified size.

``[-remaining-time <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Remaining Upgrade Time [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes where the remaining time to complete the inode
upgrade process match the specified time.

``[-scanner-progress <text>]`` - Scanner Progress [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes where the progress of the inode upgrade
process match the input.

== Examples

The following example displays information about all volumes in the middle
of the inode upgrade process on the Vserver named vs0:
```

cluster1::> volume inode-upgrade show -vserver vs0
Vserver Volume Aggregate Status %Complete Time Space Inode
Remaining Needed Progress
------- ------ --------- ------ --------- --------- ------ --------
vs0 vol1 aggr1 pending 0% - 3.07MB Public : Inode 0 out of 3822

```
:leveloffset: -1


= volume move commands

:leveloffset: +1


[[ID1c7022eaf254f2b9d51b77f1722ad853]]
= volume move abort
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop a running volume move operation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The "volume move abort" command sends an abort message to the volume move
operation and returns immediately. The volume move operation might not
abort immediately depending on the stage it is in. For example, if the
volume move operation is in a cut-over or clean-up phase, the abort is
ignored. You invoke the "volume move show" command to view the list of
running volume move operations and monitor the progress of the abort
operation. This command has the same behavior as the
xref:{relative_path}job-stop.html[job stop]`-id```_<job-id>_`` command
where the job-id is the identifier of the volume move job.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume being moved.

== Examples

The following example aborts running volume move operation on volume
``_vol1_``
```

cluster1::> volume move show
Vserver Volume State Move Phase Percent-Complete Time-To-Complete
--------- ---------- -------- ---------- ---------- ----------
vs0 vol1 alert cutover_hard_deferred 0% -
vs0 vol2 failed failed - -
2 entries were displayed.
cluster1::> volume move abort -vserver vs0 -volume vol1
cluster1::> volume move show -vserver vs0 -volume vol1 -fields completion-status
vserver volume completion-status
------- ------ ----------
vs0 vol1 "Volume move job stopped."

```
The following example shows command failed to abort on ``_vol2_`` as
volume move operation is completed.
```

cluster1::> volume move show
Vserver Volume State Move Phase Percent-Complete Time-To-Complete
--------- ---------- -------- ---------- ---------- ----------
vs0 vol1 alert cutover_hard_deferred 0% -
vs0 vol2 failed failed - -
2 entries were displayed.
cluster1::> volume move abort -vserver vs0 -volume vol2
Error: command failed: There is no volume move operation running on the
specified volume.

```
== Related Links

* xref:{relative_path}job-stop.html[job stop]




[[ID8bae1ca75e5227b935cf99fe26627c6a]]
= volume move modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify parameters for a running volume move operation

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume move modify` command modifies the parameters used by the
```

```
volume move operation in progress. These modified values can be verified
by invoking the xref:{relative_path}volume-move-show.html[volume move
show] command. The volume move operation will use the modified cutover
parameters in its next cutover attempt. Note that the modifications to the
job are not applied if the move is in the "finishing" state. This command
is not supported.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume being moved.

``[-cutover-action
{abort_on_failure|defer_on_failure|force|wait|retry_on_failure}]`` -
Specified Action For Cutover [.advaced-priv]#(privilege: advanced)#::
Specifies the action to be taken for cutover. If the effective cluster
version is ONTAP 8.3 and later, the default is ``_retry_on_failure_`` ;
otherwise the default is ``_defer_on_failure_`` . If the
``_abort_on_failure_`` action is specified, the job will try to cutover
until cutover attempts are exhausted. If it fails to cutover, it will
cleanup and end the operation. If the ``_defer_on_failure_`` action is
specified, the job will try to cutover until the cutover attempts are
exhausted. If it fails to cutover, it will move into the
"cutover_hard_deferred" state. The volume move job waits for a
xref:{relative_path}volume-move-trigger-cutover.html[volume move trigger-
cutover] command to restart the cutover process. If the ``_force_`` action
is specified, the job will try to cutover until the cutover attempts are
exhausted and force the cutover at the expense of disrupting the clients.
If the ``_wait_`` action is specified, when the job hits the decision
point, it will not go into cutover automatically, instead it will wait for
a xref:{relative_path}volume-move-trigger-cutover.html[volume move
trigger-cutover] command as the signal to try the cutover.

``[-cutover-window <integer>]`` - Specified Cutover Time Window [.advaced-
priv]#(privilege: advanced)#::
This specifies the time interval in seconds to completely cutover
operations from the original volume to the moved volume. The default value
is 30 seconds. The range for valid input is from 30 to 300 seconds,
inclusive.

``[-enforce-network-throttling {true|false}]`` - Enforce Replication
Engine Node Level Network Throttling [.advaced-priv]#(privilege:
advanced)#::
```

```
This specifies whether this move should enforce replication engine node
level network throttling. If this parameter is specified with ``_true_`` ,
the move will be throttled by the replication engine. The default is
``_false_`` .

== Examples

The following example modifies the parameters for volume move operation on
volume vol2.
```

cluster1::*> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
Volume Name: vol2
Actual Completion Time: -
Bytes Remaining: 172KB
Specified Action For Cutover: wait
Specified Cutover Time Window: 40
Time Cutover was Triggered: -
Time Cutover was last triggered: -
Destination Aggregate: node_1_aggr1
Destination Node: node_1
Detailed Status: Cutover Deferred, Waiting for user intervention(69.79MB Sent)::Volume move job preparing transfer
Estimated Time of Completion: -
Job ID: 105
Managing Node: node-2
Percentage Complete: 50%
Move Phase: cutover_hard_deferred
Estimated Remaining Duration: -
Replication Throughput: -
Duration of Move: 1 days 00:04
Source Aggregate: node_2_aggr1
Source Node: node_2
Start Time of Move: Sun Sep 18 16:40:37 2011
Move State: alert

cluster1::*> volume move modify -vserver vs0 -volume vol2 -cutover-action abort_on_failure -cutover-window 50

cluster1::*> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
Volume Name: vol2
Actual Completion Time: -
Bytes Remaining: 172KB
Specified Action For Cutover: abort_on_failure
Specified Cutover Time Window: 50
Time Cutover was Triggered: -
Time Cutover was last triggered: -
Destination Aggregate: node_1_aggr1
Destination Node: node_1
Detailed Status: Cutover Deferred, Waiting for user intervention(69.79MB Sent)::Volume move job preparing transfer

Estimated Time of Completion: -
Job ID: 106
Managing Node: node-2
Percentage Complete: 50%
Move Phase: cutover_hard_deferred
Estimated Remaining Duration: -
Replication Throughput: -
Duration of Move: 1 days 00:05
Source Aggregate: node_2_aggr1
Source Node: node_2
Start Time of Move: Sun Sep 18 16:40:37 2011
Move State: alert

```
The following example shows command failed to modify on vol1 as volume
move operation is completed.
```

cluster1::*> volume move show -vserver vs0 -volume vol1
Vserver Name: vs0
Volume Name: vol1
Actual Completion Time: Sun Sep 18 16:34:27 2011
Bytes Remaining: 172KB
Specified Action For Cutover: wait
Specified Cutover Time Window: 30
Time Cutover was Triggered: -
Time Cutover was last triggered: -
Destination Aggregate: node_1_aggr1
Destination Node: node_1
Detailed Status: Volume move failed because of a job restart
Estimated Time of Completion: -
Job ID: 108
Managing Node: node-2
Percentage Complete: -
Move Phase: failed
Estimated Remaining Duration: -
Replication Throughput: -
Duration of Move: 15 days 08:07
Source Aggregate: node_2_aggr1
Source Node: node_2
Start Time of Move: Sat Sep 03 08:27:06 2011
Move State: failed

cluster1::*> volume move modify -vserver vs0 -volume vol1 -cutover-action abort_on_failure -cutover-window 40

Error: command failed: There is no volume move operation running on the specified volume.

```
== Related Links

* xref:{relative_path}volume-move-show.html[volume move show]
* xref:{relative_path}volume-move-trigger-cutover.html[volume move
trigger-cutover]
```

```
[[ID49d80c8163e0e656510d7a6b6317752e]]
= volume move show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show status of a volume moving from one aggregate to another aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `volume move show` command displays information about volume moves in
the cluster. By default, with no parameters, it only shows volume moves
that have failed or are currently running. The command display output
depends on the parameters passed. If `-vserver` and `-volume` are
specified, the following information is displayed:

* Vserver Name: The Vserver on which the volume is located.
* Volume Name: The volume that is part of a completed or running volume
move operation.
* Actual Completion Time: The date and time in the cluster time zone when
the volume move completed.
* Bytes Remaining: The number of bytes remaining to be sent during volume
move. This is an approximate number and lags the current status by a few
minutes while the volume move is in operation.
* Specified Action for Cutover: The action to be taken for cutover or
during cutover failure. This is the input given during the start of volume
move or the value specified during a volume move modification.
* Specified Cutover Time Window: The time window in seconds given as an
input for the cutover phase of volume move. This is the input given during
the start of volume move or the value specified during a volume move
modification.
* Job ID: The Job-ID of move job.
* Destination Node: The name of the node where the destination aggregate
is present.
* Source Node: The name of the node where the source aggregate is present.
* Prior Issues Encountered: The latest issues or transient errors
encountered causing the move operation to retry the data copy phase or the
```

```
cutover phase.
* Move Initiated by Auto Balance Aggregate: The value "true" indicates the
move is initiated by Auto Balance Aggregate feature.
* Destination Aggregate: The name of the aggregate to which the volume is
moved.
* Detailed Status: The detail about any warnings, errors, and state of the
move operation.
* Estimated Time of Completion: The approximate date and time in the
cluster time zone when the entire volume move operation is expected to
complete. Note that this time may keep increasing when the move goes into
cutover-deferred mode. In those cases where the input for cutover-action
is wait, during the data copy phase, the estimated time of completion will
approximate the time to reach the cutover point and wait for user
intervention.
* Managing Node: The node in the cluster on which the move job is or was
running. This is usually on the node hosting the volume to be moved.
* Percentage Complete: The amount of work to move the volume completed
thus far in terms of percentage.
* Move Phase: The phase of the move operation.
* Estimated Remaining Duration: The approximate amount of time in terms of
days, hours, minutes and seconds remaining to complete the volume move.
* Replication Throughput: The current replication throughput of the move
operation in terms of Kb/s, Mb/s or Gb/s.
* Duration of Move: The duration in days, hours and minutes for which the
volume move was or is in progress.
* Source Aggregate: The name of the aggregate where the volume being moved
originally resides or resided.
* Start Time of Move: The date and time in the cluster time zone when the
volume move operation started.
* Move State: The state of the volume move at the time of issuing the
command and the system gathering up the information about the move.
* Original Job ID: The job-id assigned when the job was first created.
This value will only be populated when the original job-id differs from
the current job-id.



== 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.

``[-vserver <vserver name>]`` - Vserver Name::
This specifies the Vserver on which the volume is located. If this
parameter and the `-volume` parameter are specified, the command displays
detailed information about latest move performed on the specified volume.
If this parameter is specified by itself, the command displays information
about latest moves performed on volumes of the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
This specifies the volume that is part of a completed or running volume
move operation. If this parameter and the `-vserver` parameter are
specified, the command displays detailed information about latest move
performed on the specified volume. If this parameter is specified by
itself, the command displays information about the latest move on all
volumes matching the specified name.

``[-actual-completion-time <Date>]`` - Actual Completion Time::
If this parameter is specified, the command displays move operations that
match the specified date and time in the cluster time zone when the volume
move completed.

``[-bytes-remaining {<integer>[KB|MB|GB|TB|PB]}]`` - Bytes Remaining::
If this parameter is specified, the command displays move operations that
match the specified number of bytes remaining to be sent during volume
move.

``[-cutover-action
{abort_on_failure|defer_on_failure|force|wait|retry_on_failure}]`` -
Specified Action For Cutover [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified action to be taken for cutover or during cutover
failure.

``[-cutover-window <integer>]`` - Specified Cutover Time Window [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified time window in seconds for the cutover phase of volume
move.

``[-destination-aggregate <aggregate name>]`` - Destination Aggregate::
If this parameter is specified, the command displays move operations that
match the specified name of the aggregate to which the volume is being
moved.

``[-destination-node <nodename>]`` - Destination Node [.advaced-
```

```
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified name of the node where the destination aggregate is
present.

``[-details <text>]`` - Detailed Status::
If this parameter is specified, the command displays move operations that
match the specified detail about any warnings, errors and state of the
move operation.

``[-estimated-completion-time <Date>]`` - Estimated Time of Completion::
If this parameter is specified, the command displays move operations that
match the specified date and time in the cluster time zone when the entire
volume move operation is expected to complete.

``[-job-id <integer>]`` - Job ID [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified Job-ID of the move job.

``[-managing-node <nodename>]`` - Managing Node::
If this parameter is specified, the command displays move operations that
match the specified node in the cluster on which the move job is or was
running.

``[-percent-complete <percent>]`` - Percentage Complete::
If this parameter is specified, the command displays move operations that
match the specified amount of work to move the volume completed thus far
in terms of percentage.

``[-phase
{queued|initializing|replicating|cutover|cutover_hard_deferred|cutover_sof
t_deferred|aborting|completed|cleaning_up|failed|restarting|finishing}]``
- Move Phase::
If this parameter is specified, the command displays move operations that
match the specified phase of the move operation.

``[-prior-issues <text>]`` - Prior Issues Encountered [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified issues or transient errors encountered causing the
move operation to retry the data copy phase or the cutover phase.

``[-estimated-remaining-duration {<seconds>|[<d> days]
<hh>:<mm>[:<ss>]}]`` - Estimated Remaining Duration::
If this parameter is specified, the command displays move operations that
match the specified time.
```

```
``[-replication-throughput <text>]`` - Replication Throughput::
If this parameter is specified, the command displays move operations that
match the specified replication throughput of the move operation in terms
of Kb/s, Mb/s or Gb/s.

``[-actual-duration {<seconds>|[<d> days] <hh>:<mm>[:<ss>]}]`` - Duration
of Move::
If this parameter is specified, the command displays move operations that
match the specified duration for which the volume move was or is in
progress.

``[-source-aggregate <aggregate name>]`` - Source Aggregate::
If this parameter is specified, the command displays move operations that
match the specified name of the aggregate where the volume being moved
originally resides or resided.

``[-source-node <nodename>]`` - Source Node [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays move operations that
match the specified name of the node where the source aggregate is
present.

``[-start-time <Date>]`` - Start Time of Move::
If this parameter is specified, the command displays move operations that
match the specified date and time in the cluster time zone when the volume
move operation started.

``[-state {healthy|warning|alert|failed|done}]`` - Move State::
If this parameter is specified, the command displays move operations that
match the specified state of the volume move operation.

``[-moved-by-autobalance {true|false}]`` - Move Initiated by Auto Balance
Aggregate [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-original-job-id <integer>]`` - Original Job ID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-is-source-encrypted {true|false}]`` - Is Source Volume Encrypted::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-source-key-id <text>]`` - Encryption Key ID of Source Volume::
```

```
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-is-destination-encrypted {true|false}]`` - Is Destination Volume
Encrypted::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-destination-key-id <text>]`` - Encryption Key ID of Destination
Volume::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-enforce-network-throttling {true|false}]`` - Enforce Replication
Engine Node Level Network Throttling [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-is-capacity-tier-optimized {true|false}]`` - Is Capacity Tier Move
Optimized [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

``[-is-copy-free {true|false}]`` - Is Copy Free Volume Move [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays move operations that
match the specified value of this parameter.

== Examples

The following example lists status of volume move operation for a volume
vol2 on a Vserver vs0
```

cluster1::> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
Volume Name: vol2
Actual Completion Time: -
Bytes Remaining: 6.37GB
Destination Aggregate: cluster1_aggr2
Detailed Status: Transferring data: 3.67GB sent.
Estimated Time of Completion: Sat Jul 16 20:25:50 2011
Managing Node: node1
Percentage Complete: 36%
Move Phase: replicating
Estimated Remaining Duration: 00:01
Replication Throughput: 61.08MB/s
Duration of Move: 00:02

Source Aggregate: cluster1_aggr1
Start Time of Move: Sat Jul 16 20:22:01 2011
Move State: healthy

```
The following example lists status of volume move operation for a volume
vol2 on a Vserver vs0 in advanced mode
```

cluster1::*> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
Volume Name: vol2
Actual Completion Time: -
Bytes Remaining: 156KB
Specified Action For Cutover: wait
Specified Cutover Time Window: 30
Destination Aggregate: cluster1_aggr2
Destination Node: node2
Detailed Status: Cutover Deferred, Waiting for user intervention (2.04MB Sent)::Volume move job preparing transfer.
Estimated Time of Completion: -
Job ID: 265
Managing Node: node1
Percentage Complete: -
Move Phase: cutover_hard_deferred
Prior Issues Encountered: -
Estimated Remaining Duration: -
Replication Throughput: -
Duration of Move: 00:24:59
Source Aggregate: cluster1_aggr1
Source Node: node1
Start Time of Move: Tue Mar 17 22:31:32 2011
Move State: alert
Move Initiated by Auto Balance Aggregate: false
Original Job ID: -

```
The following example lists status of running and failed volume move
operations in the cluster.
```

cluster1::> volume move show
Vserver Volume State Move Phase Percent-Complete Time-To-Complete
--------- ---------- -------- ---------- ---------- ----------
vs0 s1 alert cutover_hard_deferred
98% -
vs0 vol2 failed failed - -
2 entries were displayed.

```
The following example lists status of all the volume move operations in
the cluster.
```

cluster1::> vol move show -phase *
(volume move show)
Vserver Volume State Move Phase Percent-Complete Time-To-Complete
--------- ---------- -------- ---------- --------------- ----------------
vs0 s1 alert cutover_hard_deferred
98% -
vs0 s2 done completed 100% -
vs0 vol1 failed failed - -
3 entries were displayed.

```
[[ID185786b95c15034ba9bd2671f3388be4]]
= volume move start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start moving a volume from one aggregate to another aggregate

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `volume move start` command moves a flexible volume from one storage
aggregate to another. The destination aggregate can be located on the same
node as the original aggregate or on a different node. The move occurs
within the context of the same Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the volume that will be moved.

``-destination-aggregate <aggregate name>`` - Destination Aggregate::
This specifies the aggregate to which the volume will be moved.

``[-cutover-window <integer>]`` - Cutover time window in seconds
[.advaced-priv]#(privilege: advanced)#::
This specifies the time interval to completely cutover operations from the
original volume to the moved volume. The default value is 30 seconds. The
range for valid input is from 30 to 300 seconds, inclusive.
```

``[-cutover-action {abort_on_failure|defer_on_failure|force|wait|retry_on_failure}]`` - Action for Cutover [.advaced-priv]#(privilege: advanced)#::
Specifies the action to be taken for cutover. If the effective cluster version is ONTAP 8.3 and later, the default is ``_retry_on_failure_`` ; otherwise the default is ``_defer_on_failure_`` . If the ``_abort_on_failure_`` action is specified, the job tries to cutover until cutover attempts are exhausted. If it fails to cutover, it cleans up and ends the operation. If the ``_defer_on_failure_`` action is specified, the job tries to cutover until the cutover attempts are exhausted. If it fails to cutover, it moves into the "cutover deferred" state. The volume move job waits to issue a xref:{relative_path}volume-move-trigger-cutover.html[volume move trigger-cutover] command to restart the cutover process. If the ``_force_`` action is specified, the job tries to cutover until the cutover attempts are exhausted and forces the cutover at the expense of disrupting the clients. If the ``_wait_`` action is specified, when the job hits the decision point, it does not go into cutover automatically, instead it waits to issue a xref:{relative_path}volume-move-trigger-cutover.html[volume move trigger-cutover] command as the signal to try the cutover. Once cutover is manually triggered, the cutover action changes to ``_defer_on_failure_`` . If the ``_retry_on_failure_`` action is specified, the job retries to cutover indefinitely and it never enters a "hard-deferred" state. After exhausting cutover attempts, the move job waits one hour before trying to cutover again. Issue a xref:{relative_path}volume-move-trigger-cutover.html[volume move trigger-cutover] command at any time to restart the cutover process.

``[-perform-validation-only <true>]`` - Performs validation checks only::
This is a boolean option allowing to perform pre-move validation checks for the intended volume. When set to ``_true_`` , the command only performs the checks without creating a move job. The default value is ``_false_`` .

``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the volume move operation runs as a foreground process. The default setting is ``_false_`` (that is, the operation runs in the background). Note that using this parameter will not affect how long it takes for the operation to complete.

``[-encrypt-destination {true|false}]`` - Encrypt Destination Volume::
This specifies whether the move operation should result in creating an encrypted volume on the destination aggregate. When this option is set to `true` , the destination volume will be encrypted. When it is set to `false` , the destination volume will be a plain-text volume. When this parameter is not specified, then destination will be same as the source type.

```
``[-generate-destination-key {true|false}]`` - Generate New Encryption Key
for Destination Volume::
This option is specified along with `-encrypt-destination` , a new key
will be generated, and that new key will be used for encrypting the
destination volume.

``[-tiering-policy <Tiering Policy>]`` - Volume Tiering Policy::
This optional parameter specifies the tiering policy to apply to the
destination volume. Tiering policies decide whether the user data blocks
of a volume in a FabricPool will be tiered to the cloud tier when they
become cold. FabricPool combines Flash (performance tier) with an object
store (cloud tier) into a single aggregate. The temperature of a volume
block increases if it is accessed frequently and decreases when it is not.
+
The available tiering policies are:
+
  * snapshot-only - Only the volume snapshots not associated with the
active file system are tiered to the cloud tier.
  * auto - Both snapshot data and active file system user data are tiered
to the cloud tier.
  * none - Volume blocks are not tiered to the cloud tier.
  * all - Both snapshot data and active file system user data are tiered
to the cloud tier as soon as possible without waiting for a cooling
period.



``[-allow-mixed-aggr-types {true|false}]`` - Allow Mixing FabricPool and
non-FabricPool::
If set to true, moving a FlexGroup constituent from a FabricPool to a non-
FabricPool, or vice versa, is allowed. The default value is false. This
parameter is only supported for FlexGroup constituents.

``[-encrypt-with-aggr-key {true|false}]`` - Encrypt Destination Volume
with aggr key::
This specifies whether the move operation should result in creating an
encrypted volume with aggr key on the destination aggregate. When this
option is set to `true` , the destination volume will be encrypted as NAE
(NetApp Aggregate Encryption) volume.

``[-enforce-network-throttling {true|false}]`` - Enforce Replication
Engine Node Level Network Throttling [.advaced-priv]#(privilege:
advanced)#::
This specifies whether this move should enforce replication engine node
level network throttling. If this parameter is specified with ``_true_`` ,
```

```
the move will be throttled by the replication engine. The default is
``_false_`` .

``[-is-capacity-tier-optimized {true|false}]`` - Capacity Tier Optimized
Volume Move [.advaced-priv]#(privilege: advanced)#::
This specifies whether this move should perform capacity tier optimized
volume move, in which data in the capacity tier does not need to be copied
over. If this parameter is specified as ``_true_`` and optimized volume
move is not supported, the volume move operation will fail. If this
parameter is specified as ``_false_`` , non-optimized volume move will be
performed. When this parameter is not specified, optimized volume move is
attempted, and if it is not supported, non-optimized volume move will be
performed automatically.

== Examples

The following examples perform a validation-check for a volume named
volume_test on a Vserver named vs0 to determine if it can be moved to a
destination-aggregate named dest_aggr.
```

cluster1::> volume move start -vserver vs0 -volume volume_test -destination-aggregate dest_aggr -perform -validation-only true
Error: command failed: There is 2.54GB of available space on the aggregate
dest_aggr which is not enough to accommodate a volume.
cluster1::> volume move start -vserver vs0 -volume volume_test -destination-aggregate dest_aggr -perform -validation-only true
Validation succeeded.

```
The following example performs a volume move start operation to move a
volume named volume_test on a Vserver name vs0 to a destination-aggregate
named dest_aggr.
```

cluster1::> volume move start -vserver vs0 -volume volume_test -destination-aggregate dest_aggr
[Job 267] Job is queued: Move "volume_test" in Vserver "vs0" to aggregate "dest_aggr".
Use the "volume move show -vserver vs0 -volume volume_test" command to view the status of this operation.

```
The following example performs a volume move start operation to move a
plain-text volume named vol1 to an encrypted volume on destination-
aggregate aggr1.
```

cluster1::> volume move start -volume vol1 -destination-aggregate aggr1 -encrypt-destination true
[Job 267] Job is queued: Move "vol1" in Vserver "vs1" to aggregate "aggr1".
Use the "volume move show -vserver vs1 -volume vol1" command to view the status of this operation.

```
== Related Links
```

```
* xref:{relative_path}volume-move-trigger-cutover.html[volume move
trigger-cutover]




[[ID2e04fc68a4dbddaed80d49f70818e706]]
= volume move trigger-cutover
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Trigger cutover of a move job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command causes a replicating or deferred volume move job to attempt
cutover. Unless the force option is set, cutover entry is not guaranteed.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
The Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
The volume that is being moved.

``[-force <true>]`` - Force Cutover::
If this parameter is specified, the cutover is done without confirming the
operation - even if the operation could cause client I/O disruptions.

== Examples

cluster1::> volume move trigger-cutover -vserver vs0 -volume testvol_1
-force




[[ID8fd5e36c103a03eb5492a1f062a0b18f]]
= volume move recommend show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Move Recommendations

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume move recommend show` command displays moves that were
recommended by the Auto Balance Aggregate feature.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the display will be limited to only those
recommendations with a Vserver that matches the specified value.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the display will be limited to only those
recommendations with a volume that matches the specified value.

``[-creation-time <MM/DD/YYYY HH:MM:SS>]`` - Time Stamp of Recommendation
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the display will be limited to only those
recommendations with a creation-time that matches the specified value.

``[-source-aggregate <aggregate name>]`` - Unbalanced Aggregate Name
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the display will be limited to only those
recommendations with a source-aggregate that matches the specified value.
```

```
``[-source-space-after <percent>]`` - Space Free After Move (%) [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the display will be limited to only those
recommendations with a source-space-after that matches the specified
value.

``[-destination-aggregate <aggregate name>]`` - Destination Aggregate Name
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the display will be limited to only those
recommendations with a destination-aggregate that matches the specified
value.

``[-destination-space-after <percent>]`` - Space Bump After Move (%)
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the display will be limited to only those
recommendations with a destination-space-after that matches the specified
value.

== Examples

The following example displays information about the recommendations made
by the Auto Balance Aggregate feature.
```

cluster1::*> volume move recommend show -instance
Vserver Name: vs0.example.com
Volume Name: ro10
Time Stamp of Recommendation: 3/13/2014 16:26:39
Unbalanced Aggregate Name: aggr_1
Space Free After Move (%): 36%
Destination Aggregate Name: aggr_3
Space Bump After Move (%): 36%

```
[[ID3a5fb1584863e1a5531248acedfdf5a6]]
= volume move target-aggr show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
List target aggregates compatible for volume move

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description
```

```
 The volume move target-aggr show displays information about compatible
target aggregates for the specified volume to be moved to.

== 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.

``[-vserver <vserver name>]`` - Vserver Name (Required field)::
Selects information about compatible target aggregates for volumes of the
specified Vserver.

``[-volume <volume name>]`` - Volume Name (Required field)::
Selects information about compatible target aggregates that have enough
space for the specified volume.

``[-aggregate <aggregate name>]`` - Aggregate Name::
Selects information about compatible target aggregates with the specified
aggregate name (to which the volume might be moved).

``[-tiering-policy <Tiering Policy>]`` - Volume Tiering Policy::
Selects information about compatible target aggregates with the specified
destination tiering policy.

``[-availsize {<integer>[KB|MB|GB|TB|PB]}]`` - Available size::
Selects information about compatible target aggregates that have the
specified available size.

``[-storagetype <text>]`` - Storage Type::
Selects information about compatible target aggregates with the specified
storage type. Examples of storage types are “ATA”, ”BSAS”, “FCAL”, “LUN”,
“SATA”, “SAS” and “SSD”.

``[-allow-mixed-aggr-types {true|false}]`` - Allow Mixing FabricPool and
non-FabricPool::
If set to true, moving a FlexGroup constituent from a FabricPool to a non-
FabricPool, or vice versa, is allowed. The default value is false. This
parameter is only supported for FlexGroup constituents.

== Examples
```

```
The following example lists target aggregates compatible for moving a
volume vol1 on a Vserver vs1.
```

cluster1::> volume move target-aggr show -vserver vs1 -volume vol1
Aggregate Name Available Size Storage Type
--- --- ---
aggr1 113.5GB FCAL
aggr2 113.5GB FCAL
2 entries were displayed.

```
:leveloffset: -1


= volume object-store commands

:leveloffset: +1


[[ID645e8c6c014400f358da287f72aa414c]]
= volume object-store tiering show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Tiering Status of FabricPool Volumes

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume object-store tiering show` command displays information about
the tiering scanner status of FabricPool aggregate volumes. The tiering
scanner will, depending on the combination of the tiering-policy and
cloud-retrieval-policy settings, both retrieve from and tier to the
capacity tier. This show command can be used to display the status of the
scanner, whether it has aborted, what errors it has encountered and when
it will be scheduled to run again.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
vserver name

``[-volume <volume name>]`` - Volume [.advaced-priv]#(privilege:
advanced)#::
volume name

``[-node <nodename>]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
node name

``[-vol-dsid <integer>]`` - Volume DSID [.advaced-priv]#(privilege:
advanced)#::
volume dsid

``[-aggregate <aggregate name>]`` - Aggregate Name [.advaced-
priv]#(privilege: advanced)#::
aggregate name

``[-scanner-state {ready|active|waiting|waiting_work}]`` - State
[.advaced-priv]#(privilege: advanced)#::
The state of the scanner. The values active, ready and waiting may appear
in this field. Active means the scanner is currently running, waiting
means the scanner is waiting to be scheduled, and ready means the scanner
is ready to run. Once a scanner is ready to run, it must be scheduled by
the system's scan scheduler, the amount of time it takes for it to become
active is variable and depends on how busy the system is..

``[-scanner-last-status {aborted-nospace-promote|aborted-
fabricpool|aborted-nospace-defrag|aborted-nospace|aborted-
exception|aborted-policy|aborted|completed|prev-not-run|aborted-defrag-
throttled|aborted-nospace-revert-format|aborted-revert-format-
paused|aborted-dual-tier-dirty-failed}]`` - Previous Run Status [.advaced-
priv]#(privilege: advanced)#::
The status of the tiering scanner the last time it ran. This field will
indicate if the scanner has aborted or completed. Possible status options
are:
+
  * aborted-nospace-promote : The tiering scanner was aborted because of
```

```
an internal space check error or a lack of space while promoting data from
the capacity tier to the performance tier. Check the promotion-space-
needed field.
  * aborted-fabricpool: The tiering scanner was aborted because it tried
running inactive data reporting on a FabricPool volume.
  * aborted-nospace-defrag : The tiering scanner was aborted because of an
internal space check error while defragmenting.
  * aborted-nospace : The tiering scanner was aborted because of an
internal space check error.
  * aborted-exception : The tiering scanner encountered an exception while
reading data from the capacity tier. Check the scanner-abort-exception
field to see the error code that caused the exception.
  * aborted-policy : The tiering scanner was aborted because either the
tiering-policy or the cloud-retrieval-policy changed while the scan was in
progress.
  * aborted : The tiering scan was manually aborted.
  * completed : The tiering scanner was completed successfully.
  * prev-not-run : The tiering scanner has not run before.
  * aborted-defrag-throttled : The tiering scanner was aborted because the
defragmentation cannot continue due to its threshold limit.
  * aborted-nospace-revert-format : The tiering scanner was aborted
because of an internal space check error while reverting unified object
format to native format.
  * aborted-revert-format-paused: The tiering scanner was aborted because
it failed to drain waiting objects while reverting unified object format
to native format.
  * aborted-dual-tier-dirty-failed: The tiering scanner was aborted
because it failed to tier blocks while reverting unified object format to
native format.



``[-scanner-abort-exception <integer>]`` - Aborted Exception Status
[.advaced-priv]#(privilege: advanced)#::
If the scanner last status indicates the scanner aborted with an exception
then this field is set to the exception error code that caused the abort.

``[-time-last-scan <Date>]`` - Time Scanner Last Finished [.advaced-
priv]#(privilege: advanced)#::
This field displays the date and time the last scanner status field was
set.

``[-scan-percent <percent>]`` - Scanner Percent Complete [.advaced-
priv]#(privilege: advanced)#::
This field displays the completion percentage of scans in the active
state.
```

```
``[-time-next-scan <Date>]`` - Time Waiting Scan will be scheduled
[.advaced-priv]#(privilege: advanced)#::
This field displays the date and time scans in the waiting state will
enter the ready state.

``[-tiering-policy <Tiering Policy>]`` - Tiering Policy [.advaced-
priv]#(privilege: advanced)#::
The tiering policy. This field displays the current tiering-policy in
effect on the volume.

``[-promotion-space-needed {<integer>[KB|MB|GB|TB|PB]}]`` - Estimated
Space Needed for Promotion [.advaced-priv]#(privilege: advanced)#::
If the scanner aborted due to no space for promote, then this field is set
to the estimated minimum space required for promotion to take effect.

``[-scan-start-time <Date>]`` - Time Scan Started [.advaced-
priv]#(privilege: advanced)#::
This field displays the date and time the scan, if in active state,
started.

``[-cloud-retrieval-policy {default|on-read|never|promote}]`` - Cloud
Retrieval Policy [.advaced-priv]#(privilege: advanced)#::
The cloud retrieval policy. This field displays the current cloud-
retrieval-policy in effect on the volume.

``[-scan-elapsed-time
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Elapsed Time
Scanner Ran [.advaced-priv]#(privilege: advanced)#::
This field displays the time it took for the last scan to run to complete
or abort.

``[-revert-scan-state <text>]`` - Object Format Revert Scan State
[.advaced-priv]#(privilege: advanced)#::
This field displays the state of object format revert scan. The values
inactive, ready, active and complete may appear in this field. Inactive
means revert scan is not running on this volume, ready means scan is
waiting to start running, active mean scanner is running and complete
indicates revert scan is complete on this volume.

``[-scan-last-space-err <integer>]`` - space error code encountered in
last scan [.advaced-priv]#(privilege: advanced)#::
This field displays the most recent space check error code, if any, that
was encountered.

``[-blocks-promoted <integer>]`` - Blocks Promoted [.advaced-
priv]#(privilege: advanced)#::
```

```
This field displays the number of blocks that were promoted from the
capacity tier to the performance tier.

== Examples
```

cluster1::> volume object-store tiering show
Policy
Vserver Volume Aggregate State %Complete Tiering CloudRetrieve LastStatus
------- ------ --------- -------- --------- --------- ------------- ----------
vs1 vol aggr1 waiting - snapshot-only
default completed

```
Shows the tiering status for all FabricPool volumes.




[[ID63d2dc4221ba90d73bcc1939bc5deb17]]
= volume object-store tiering trigger
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Trigger a tiering scan

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume object-store tiering trigger` command triggers a tiering and
retrieve scan. Tiering and retrieval behavior is driven by the tiering and
cloud retrieve policy settings on the volume. The cloud retrieve policy
must be set to promote to enable scanner based retrieval.

== Parameters

``-vserver <vserver name>`` - VServer Name [.advaced-priv]#(privilege:
advanced)#::
vserver name

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
volume name
```

```
:leveloffset: -1


= volume qtree commands

:leveloffset: +1


[[ID6de281799c767af220976b2252d6a318]]
= volume qtree create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new qtree

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates a qtree in the Vserver and volume you specify. You
can create up to 4,994 qtrees per volume.

You can optionally specify the following attributes when creating a new
qtree:



* Security style
* Opportunistic lock mode
* User ID
* Group ID
* UNIX permissions
* Export Policy



== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume containing the
```

```
qtree belongs.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume that will contain the qtree you are
creating.

``-qtree <qtree name>`` - Qtree Name::
This specifies the name of the qtree you are creating.
+
A qtree name cannot contain a forward slash (/). The qtree name cannot be
more than 64 characters long.

| ``-qtree-path <qtree path>`` - Actual (Non-Junction) Qtree Path }::
The qtree path argument in the format /vol/<volume name>/<qtree name> can
be specified instead of specifying volume and qtree as separate arguments.

``[-security-style <security style>]`` - Security Style::
This optionally specifies the security style for the qtree, which
determines how access to the qtree is controlled. The supported values are
``_unix_`` (for UNIX uid, gid and mode bits), ``_ntfs_`` (for CIFS ACLs),
and ``_mixed_`` (for NFS and CIFS access). If you do not specify a
security style for the qtree, it inherits the security style of its
containing volume.

``[-oplock-mode {enable|disable}]`` - Oplock Mode::
This optionally specifies whether oplocks are enabled for the qtree. If
you do not specify a value for this parameter, it inherits the oplock mode
of its containing volume.

``[-user <user name>]`` - User ID::
This optionally specifies the name or ID of the user that is set as the
owner of the qtree.

``[-group <group name>]`` - Group ID::
This optionally specifies the name or ID of the group that is set as the
owner of the qtree.

``[-m, -unix-permissions <unix perm>]`` - Unix Permissions::
This optionally specifies the UNIX permissions for the qtree when the `-
security-style` is set to `unix` or `mixed.` You can specify UNIX
permissions either as a four-digit octal value (for example, 0700) or in
the style of the UNIX ls command (for example, ` -rwxr-x---` ). For
information on UNIX permissions, see the UNIX or Linux documentation. If
you do not specify UNIX permissions for the qtree, it inherits the UNIX
permissions of its containing volume.
```

```
``[-export-policy <text>]`` - Export Policy::
This optional parameter specifies the name of the export policy associated
with the qtree. For information on export policies, see the documentation
for the xref:{relative_path}vserver-export-policy-create.html[vserver
export-policy create] command. If you do not specify a value for this
parameter, it inherits the export policy of its containing volume.

``[-qos-policy-group <text>]`` - QoS policy group::
A policy group defines measurable service level objectives (SLOs) that
apply to the storage objects with which the policy group is associated. If
you do not assign a policy group to a qtree, the system does not monitor
and control the traffic to it.

== Examples

The following example creates a qtree named qtree1. The Vserver name is
vs0 and the volume containing the qtree is named vol1. The qtree has a
mixed security style. Its other attributes are inherited from volume vol1.
```

cluster1::> volume qtree create -vserver vs0 -volume vol1 -qtree qtree1 -security-style mixed

```
The following example uses a 7G-compatible command to create the qtree.
```

cluster1::> vserver context vs0
vs0::> qtree create /vol/vol1/qtree1

```
== Related Links

* xref:{relative_path}vserver-export-policy-create.html[vserver export-
policy create]




[[ID6b1c928a400b6045a68b97919fdbe2cb]]
= volume qtree delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a qtree

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

This command deletes a qtree. The length of time that it takes to delete a
qtree depends on the number of directories and files it contains. You can
monitor the progress of the delete operation by using the
xref:{relative_path}job-show.html[job show] and xref:{relative_path}job-
watch-progress.html[job watch-progress] commands, respectively.

The automatically created qtree in the volume - qtree0, listed in CLI
output as "" - cannot be deleted.

NOTE: Quota rules associated with this qtree in all the quota policies
will be deleted when you delete this qtree. Qtree deletion will not be
allowed if Storage-level Access Guard (SLAG) is configured.



== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume containing the
qtree belongs.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume containing the qtree to be deleted.

``-qtree <qtree name>`` - Qtree Name::
This specifies the name of the qtree to be deleted.

| ``-qtree-path <qtree path>`` - Actual (Non-Junction) Qtree Path }::
The qtree path argument in the format /vol/<volume name>/<qtree name> can
be specified instead of specifying volume and qtree as separate arguments.

``[-force <true>]`` - Force Delete [.advaced-priv]#(privilege:
advanced)#::
This optionally forces the qtree delete operation to proceed when the
qtree contains files. The default setting is false (that is, the qtree
will not be deleted if it contains files). This parameter is available
only at the advanced privilege and higher.

``[-foreground <true>]`` - Foreground Process::
This optionally specifies whether the qtree delete operation runs as a
foreground process. The default setting is false (that is, the operation
runs in the background).
```

```
== Examples

The following example deletes a qtree named qtree4. The Vserver name is
vs0 and the volume containing the qtree is named vol1.
```

cluster1::> volume qtree delete -vserver vs0 -volume vol1 -qtree qtree4
WARNING: Are you sure you want to delete qtree qtree4 in volume vol1 vserver vs0? {y|n}: y
[Job 38] Job is queued: Delete qtree qtree4 in volume vol1 vserver vs0.

```
== Related Links

* xref:{relative_path}job-show.html[job show]
* xref:{relative_path}job-watch-progress.html[job watch-progress]




[[IDf69b51ea0ce8bfdac38f8549571ba822]]
= volume qtree modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify qtree attributes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command allows you to modify the following attributes of an existing
qtree in the given Vserver and volume:

* Security style
* Opportunistic lock mode
* User ID
* Group ID
* UNIX permissions
* Export policy



== Parameters
```

```
``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume containing the
qtree belongs.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume containing the qtree to be modified.

``-qtree <qtree name>`` - Qtree Name::
This specifies the name of the qtree to be modified. You can modify the
attributes of qtree0 (represented as "" in the CLI) by omitting the `-
qtree` parameter from the command or by specifying the value `""""` for
the `-qtree` parameter.

| ``-qtree-path <qtree path>`` - Actual (Non-Junction) Qtree Path }::
The qtree path argument in the format /vol/<volume name>/<qtree name> can
be specified instead of specifying volume and qtree as separate arguments.
The automatically created qtree0 can be represented as /vol/<volume name>.

``[-security-style <security style>]`` - Security Style::
This optionally modifies the security style for the qtree. The supported
values are ``_unix_`` (for UNIX uid, gid and mode bits), ``_ntfs_`` (for
CIFS ACLs), and ``_mixed_`` (for NFS and CIFS access). Modifying a qtree's
security style will not affect any of the files in the other qtrees of
this volume.

``[-oplock-mode {enable|disable}]`` - Oplock Mode::
This optionally modifies whether oplocks are enabled for the qtree.
+
Modifying qtree0's oplock mode will not affect any of the files in the
other qtrees of this volume.

``[-user <user name>]`` - User ID::
This optionally modifies the name or ID of the user that is set as the
owner of the qtree.

``[-group <group name>]`` - Group ID::
This optionally modifies the name or ID of the group that is set as the
owner of the qtree.

``[-unix-permissions <unix perm>]`` - Unix Permissions::
This optionally modifies the UNIX permissions for the qtree. You can
specify UNIX permissions either as a four-digit octal value (for example,
0700) or in the style of the UNIX ls command (for example, ` -rwxr-x---`
). For information on UNIX permissions, see the UNIX or Linux
documentation.
+
```

```
The unix permissions can be modified only for qtrees with unix or mixed
security style.

``[-export-policy <text>]`` - Export Policy::
This optional parameter modifies the export policy associated with the
qtree. If you do not specify an export policy name, the qtree inherits the
export policy of the containing volume.For information on export policy,
see the documentation for the xref:{relative_path}vserver-export-policy-
create.html[vserver export-policy create] command.

``[-qos-policy-group <text>]`` - QoS Policy Group::
This optional parameter specifies which QoS policy group to apply to the
qtree. This policy group defines measurable service level objectives
(SLOs) that apply to the storage objects with which the policy group is
associated. If you do not assign a policy group to a qtree, the system
does not monitor and control the traffic to it. To remove this qtree from
a policy group, enter the reserved keyword "none".

== Examples

The following example modifies a qtree named qtree1. The Vserver name is
vs0 and the volume containing the qtree is named vol1. The qtree now has a
UNIX security style and oplocks are enabled.
```

cluster1::> volume qtree modify -vserver vs0 -volume vol1 -qtree qtree1 -security-style unix -oplocks enabled

```
== Related Links

* xref:{relative_path}vserver-export-policy-create.html[vserver export-
policy create]




[[IDb83777723f6346c05e4dc7acb51008c7]]
= volume qtree oplocks
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify qtree oplock mode

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

This command allows you to display or modify the opportunistic lock mode
of a qtree.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume containing the
qtree belongs.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume containing the qtree.

``-qtree <qtree name>`` - Qtree Name::
This specifies the name of the qtree for which the oplock mode is being
displayed or modified.

| ``-qtree-path <qtree path>`` - Actual (Non-Junction) Qtree Path }::
The qtree path argument in the format /vol/<volume name>/<qtree name> can
be specified instead of specifying volume and qtree as separate arguments.
The automatically created qtree0 can be represented as /vol/<volume name>.

``[-oplock-mode {enable|disable}]`` - Oplock Mode::
This specifies the new oplock mode of the qtree. If this parameter is not
specified, then the current oplock mode of the qtree is displayed.
+
Modifying qtree0's oplock mode will not affect any of the files in the
other qtrees of this volume.

== Examples

The following example displays the oplock mode of a qtree called qtree1.
The Vserver name is vs0 and the volume containing the qtree is named vol1.
```

cluster1::> volume qtree oplocks -vserver vs0 -volume vol1 -qtree qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.

```
 The following example modifies the oplock mode of a qtree called qtree2
to enabled. The Vserver name is vs0 and the volume containing the qtree is
named vol1.
```

cluster1::> volume qtree oplocks -vserver vs0 -volume vol1 -qtree qtree2 -oplock-mode enable

```
The following example uses a 7G-compatible command to display and modify
the oplock mode of a qtree.
```

cluster1::> vserver context vs0
vs0::> qtree oplocks /vol/vol1/qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.
vs0::> qtree oplocks /vol/vol1/qtree2 enable

```
[[IDf82874df3e87eabe8f43af66f2c4a855]]
= volume qtree rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename an existing qtree

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command allows you to rename an existing qtree.

The automatically created qtree in the volume - qtree0, listed in CLI
output as "" - cannot be renamed.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume containing the
qtree belongs.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume containing the qtree to be renamed.

``-qtree <qtree name>`` - Qtree Name::
This specifies the name of the qtree to be renamed.

| ``-qtree-path <qtree path>`` - Actual (Non-Junction) Qtree Path }::
The qtree path argument in the format /vol/<volume name>/<qtree name> can
be specified instead of specifying volume and qtree as separate arguments.

``[-newname <qtree name>]`` - Qtree New Name::
This specifies the new name of the qtree. The new qtree name cannot
contain a forward slash (/) and cannot be more than 64 characters long.

== Examples

The following example renames a qtree named qtree3 to qtree4. The Vserver
name is vs0 and the volume containing the qtree is named vol1.
```

cluster1::> volume qtree rename -vserver vs0 -volume vol1 -qtree qtree3 -newname qtree4

```
[[ID60b310f888043199f7a435001b53daee]]
= volume qtree security
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify qtree security style

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command allows you to display or modify the security style of a
qtree.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume containing the
qtree belongs.

{ ``-volume <volume name>`` - Volume Name::
This specifies the name of the volume containing the qtree.

``-qtree <qtree name>`` - Qtree Name::
This specifies the name of the qtree for which the security style is being
displayed or modified.

| ``-qtree-path <qtree path>`` - Actual (Non-Junction) Qtree Path }::
The qtree path argument in the format /vol/<volume name>/<qtree name> can
be specified instead of specifying volume and qtree as separate arguments.
The automatically created qtree0 can be represented as /vol/<volume name>.

``[-security-style <security style>]`` - Security Style::
This specifies the new security style of the qtree. If this parameter is
not specified, then the current security style of the qtree is displayed.
The supported values are ``_unix_`` (for UNIX uid, gid and mode bits),
``_ntfs_`` (for CIFS ACLs), and ``_mixed_`` (for NFS and CIFS access).
Modifying a qtree's security style will not affect any of the files in the
other qtrees of this volume.
```

```
== Examples

The following example displays the security style of a qtree called
qtree1. The Vserver name is vs0 and the volume containing the qtree is
named vol1.
```

cluster1::> volume qtree security -vserver vs0 -volume vol1 -qtree qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.

```
 The following example modifies the security style of a qtree called
qtree2 to unix. The Vserver name is vs0 and the volume containing the
qtree is named vol1.
```

cluster1::> volume qtree security -vserver vs0 -volume vol1 -qtree qtree2 -security-style unix

```
The following example uses a 7G-compatible command to display and modify
the security style of a qtree.
```

cluster1::> vserver context vs0
vs0::> qtree security /vol/vol1/qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.
vs0::> qtree security /vol/vol1/qtree2 unix

```
[[IDddf71d867ea10addfc2ebab2cd9415a4]]
= volume qtree show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of qtrees

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays information about qtrees for online volumes. By
default, the command displays the following information about all qtrees
in the cluster:

* Vserver name
* Volume name
```

```
* Qtree name
* Security style (unix, ntfs, mixed or unified)
* Whether oplocks is enabled
* Status (normal or readonly)



The display will also include information about Qtree 0. When you create a
volume, a special qtree referred to as "qtree0", also called the default
qtree is automatically created for the volume. It represents all of the
data stored in a volume that is not contained in a qtree. In the CLI
output, qtree0 is denoted by empty quotation marks ("") and has the ID
zero (0). The qtree called qtree0 cannot be manually created or deleted.

The qtree status indicates readonly for data protection and load sharing
volumes.

To display detailed information about a single qtree, run the command with
the -instance and -qtree parameters. The detailed view adds the following
information:

* User ID
* Group ID
* UNIX permissions
* Qtree ID
* Export policy
* Is Export Policy Inherited



== 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.

| ``[-exports  ]``::
Displays the following information about qtree exports:
+
  * Vserver - The name of the Vserver the qtree belongs to
  * Volume - The name of the volume the qtree resides on
  * Qtree name - The name of the qtree
  * Policy Name - The name of the export policy assigned to the qtree
  * Is Export Policy Inherited - Whether the export policy assigned to the
qtree is inherited
```

```
| ``[-id  ]``::
Displays qtree IDs in addition to the default output.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name::
Selects information about the qtrees in the specified Vserver.

{ ``[-volume <volume name>]`` - Volume Name::
Selects information about the qtrees in the specified volume.

``[-qtree <qtree name>]`` - Qtree Name::
Selects information about the qtrees that have the specified name.

| ``[-qtree-path <qtree path>]`` - Actual (Non-Junction) Qtree Path }::
Selects information about the qtrees that have the specified path.

``[-security-style <security style>]`` - Security Style::
Selects information about the qtrees that have the specified security
style.

``[-oplock-mode {enable|disable}]`` - Oplock Mode::
Selects information about the qtrees that have the specified oplock mode.

``[-user <user name>]`` - User ID::
Selects information about the qtrees that have the specified user.

``[-group <group name>]`` - Group ID::
Selects information about the qtrees that have the specified group.

``[-m, -unix-permissions <unix perm>]`` - Unix Permissions::
Selects information about the qtrees that have the specified UNIX
permissions.

``[-qtree-id <integer>]`` - Qtree Id::
Selects information about the qtrees that have the specified ID. A valid
qtree ID is an integer from 0 to 4994. All qtree0 (automatically created)
qtrees have an ID of zero (0).

``[-status {normal|readonly}]`` - Qtree Status::
Selects information about the qtrees that have the specified status.

``[-export-policy <text>]`` - Export Policy::
```

```
Selects information about the qtrees that use the specified export policy.

``[-is-export-policy-inherited {true|false}]`` - Is Export Policy
Inherited::
Selects information about the qtrees that inherit (true) or not inherit
(false) the export policy of containing volume.

``[-qos-policy-group <text>]`` - QoS policy group::
Selects information about the qtrees that use the specified QoS policy.

== Examples

The following example displays default information about all qtrees along
with each qtree ID. Note that on vs0, no qtrees have been manually
created, so only the automatically created qtrees referred to as qtree 0
are shown. On vs1, the volume named vs1_vol1 contains qtree 0 and two
manually created qtrees, qtree1 and qtree2.
```

cluster1::> volume qtree show -id
Vserver Volume Qtree Style Oplocks Status Id
--- --- --- --- --- --- — vs0 vs0_vol1 "" unix enable readonly 0
vs0 vs0_vol2 "" unix enable normal 0
vs0 vs0_vol3 "" unix enable readonly 0
vs0 vs0_vol4 "" unix enable readonly 0
vs0 root_vs_vs0 "" unix enable normal 0
vs1 vs1_vol1 "" unix enable normal 0
vs1 vs1_vol1 qtree1 unix disable normal 1
vs1 vsl_vol1 qtree2 unix enable normal 2
vs1 root_vs_vs1 "" unix enable normal 0
9 entries were displayed.

```
:leveloffset: -1


= volume quota commands

:leveloffset: +1


[[IDa323b91dcac23e96d30d4b9b7956fed3]]
= volume quota modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Modify quota state for volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command allows you to modify the following quota attributes for one
or more volumes:

* Quota state
* Whether quota exceeded messages are logged or not
* Frequency with which quota exceeded messages are logged



Modifications to the quota state for a volume creates a job to perform the
quota state changes for that volume. You can monitor the progress of the
job by using the xref:{relative_path}job-show.html[job show] and
xref:{relative_path}job-watch-progress.html[job watch-progress] commands.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume whose quota
attributes you are modifying is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume whose quota attributes you are
modifying.

``[-state <quota_state>]`` - Quota State::
This parameter optionally modifies the quota state to one of the
following:
+
  * `off` - This indicates that quotas be deactivated for the specified
volume.
  * `on` - This indicates that quotas be activated for the specified
volume.
  * `resize` - This indicates that the quota limits be resized according
to the values specified in the quota policy assigned to the Vserver. Note
that quotas must be activated first for a volume before a resize operation
can be performed.


+
Both quota activation and quota resize operations apply the quota rules
```

```
configured for the volume within the quota policy that is currently
assigned to the Vserver. These quota rules are managed by using the
commands in the `volume quota policy rule` menu. Quotas, when activated
for a volume, go through an initialization process. As part of the quota
initialization all the quota rules are applied to the volume. In addition,
a filesystem scanner is started to scan the entire filesystem within the
volume to bring the quota accounting and reporting up to date. The quota
job finishes after the filesystem scanner is started on the volume. The
quota state for the volume is `initializing` until the filesystem scanner
finishes scanning the entire filesystem. After the scanning is complete,
the quota state will be `on` .
+
When quotas are resized, the quota state is `resizing` until the resizing
operation finishes. As part of this operation, the quota limits for quotas
currently in effect are resized to the limits currently configured for the
volume. After the quota resize operation finishes, the quota state will be
`on` .
+
Quota state changes can also be performed using the commands
xref:{relative_path}volume-quota-on.html[volume quota on] ,
xref:{relative_path}volume-quota-off.html[volume quota off] and
xref:{relative_path}volume-quota-resize.html[volume quota resize] .

``[-logging {on|off}]`` - Logging Messages::
This parameter optionally specifies whether quota exceeded syslog/EMS
messages are logged in the system log messages. When it is set to `on` ,
quota exceeded messages are generated when the user exceeds the quota's
disk limit or the file limit through a NFS/CIFS operation or any operation
within the ONTAP software. When set to `off` no quota exceeded messages
are generated. This parameter can be changed only after quotas are
activated for a volume.

``[-logging-interval <text>]`` - Logging Interval::
This parameter optionally specifies a logging interval, which indicates
the frequency with which quota exceeded messages are generated. You can
specify a logging interval in the `<integer><suffix>` format, where suffix
can be minutes (`m` ), hours (`h` ), or days (`d` ), but not combinations
thereof (in other words, `90m` is a valid logging interval, but `1h30m` is
not a valid logging interval). You can modify the logging interval only
when the logging is `on` . When quotas are first activated, the logging
parameter is automatically set to `on` , and the logging interval set to
`1h` . If continuous logging is required, an interval of `0m` should be
specified. This parameter can be changed only after quotas are activated
for a volume.
+
NOTE: quota message logging may not occur at exactly the same interval
```

```
rate as specified by the user, especially for very small intervals. This
is due to the behavior of the logging system that buffers messages instead
of outputting them immediately. Setting the logging interval to `0m` can
cause lots of quota exceeded messages to be logged in the system log
messages.
+


``[-foreground <true>]`` - Foreground Process::
This parameter optionally specifies whether the job created by quota state
modify operation runs as a foreground process. The default setting is
`false` (that is, the quota state modify operation runs in the
background). When set to `true` , the command will not return until the
job completes.

== Examples

The following example activates quotas on the volume named `vol1` , which
exists on Vserver `vs0` .
```

cluster1::> volume quota modify -vserver vs0 -volume vol1 -state on
[Job 24] Job is queued: Quota ON Operation on vserver vs0 volume vol1.

```
The following example turns on quota message logging and sets the logging
interval to 4 hours.
```

cluster1::> volume quota modify -vserver vs0 -volume vol1 -logging on -logging-interval 4h

```
The following example resizes quota limits on a volume.
```

cluster1::> volume quota modify -vserver vs0 -volume vol1 -state resize -foreground true
[Job 80] Job succeeded: Successful

```
== Related Links

* xref:{relative_path}job-show.html[job show]
* xref:{relative_path}job-watch-progress.html[job watch-progress]
* xref:{relative_path}volume-quota-on.html[volume quota on]
* xref:{relative_path}volume-quota-off.html[volume quota off]
* xref:{relative_path}volume-quota-resize.html[volume quota resize]
```

```
[[IDe9780e3b4b8b5cd3a506f58a0907a74d]]
= volume quota off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Turn off quotas for volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates a job to deactivate quotas for the specified volume.
You can monitor the progress of the job by using the
xref:{relative_path}job-show.html[job show] and xref:{relative_path}job-
watch-progress.html[job watch-progress] commands.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which you are deactivating
quotas.

``[-foreground <true>]`` - Foreground Process::
This optionally specifies whether the job created for deactivating quotas
runs as a foreground process. The default setting is `false` (that is, the
operation runs in the background). When set to `true` , the command will
not return until the job completes.

== Examples

The following example deactivates quotas on the volume named `vol1` ,
which exists on Vserver `vs0` .
```

cluster1::> volume quota off -vserver vs0 -volume vol1
[Job 23] Job is queued: Quota OFF Operation on vserver vs0 volume vol1.

```
The following example uses a 7G-compatible command to deactivate quotas on
the volume named `vol1` which exists on Vserver `vs0` .
```

cluster1::> vserver context vs0
vs0::> quota off vol1
[Job 25] Job is queued: Quota OFF Operation on vserver vs0 volume vol1.

```
== Related Links

* xref:{relative_path}job-show.html[job show]
* xref:{relative_path}job-watch-progress.html[job watch-progress]




[[ID7454aadbf29558c94b5676f2762dab3d]]
= volume quota on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Turn on quotas for volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates a job to activate quotas for the specified volume.
You can monitor the progress of the job by using the
xref:{relative_path}job-show.html[job show] and xref:{relative_path}job-
watch-progress.html[job watch-progress] commands.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which you are activating quotas.

``[-w, -foreground <true>]`` - Foreground Process::
This optionally specifies whether the job created for activating quotas
runs as a foreground process. The default setting is `false` (that is, the
operation runs in the background). When set to `true` , the command will
not return until the job completes. The quota job finishes after the
filesystem scanner is started. The quota state for the volume is
```

```
`initializing` until the filesystem scanner finishes scanning the entire
filesystem. After the scanning is complete, the quota state will be `on` .

== Examples

The following example activates quotas on the volume named `vol1` , which
exists on Vserver `vs0` .
```

cluster1::> volume quota on -vserver vs0 -volume vol1
[Job 23] Job is queued: Quota ON Operation on vserver vs0 volume vol1.

```
The following example uses a 7G-compatible command to activate quotas on
the volume named `vol1` which exists on Vserver `vs0` .
```

cluster1::> vserver context vs0
vs0::> quota on -w vol1
[Job 25] Job is queued: Quota ON Operation on vserver vs0 volume vol1.

[Job 25] Job succeeded: Successful

```
== Related Links

* xref:{relative_path}job-show.html[job show]
* xref:{relative_path}job-watch-progress.html[job watch-progress]




[[IDc8f8514eac4494b38bd1f6dc8dc46d26]]
= volume quota report
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the quota report for volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays the quota report for all volumes in each Vserver
that are online and for which quotas are activated. Quota report includes
```

```
the quota rules (default, explicit, and derived) in effect and the
associated resource usage (disk space and files). If quotas are still
initializing for a specific volume, that volume is not included.

This command displays the following information:

* Vserver name
* Volume name
* Index - This is a unique number within a volume assigned to each quota
rule displayed in the quota report.
* Tree name - This field gives the name of the qtree if the quota rule is
at the qtree level. It is empty if the quota rule is at the volume level.
* Quota type - Type of quota rule (`tree` or `user` or `group` ).
* Quota target - This field gives the name of the target of the quota
rule. For tree quota rules, it will be the qtree ID of the qtree. For user
quota rules, it will be the UNIX user name or the Windows user name. For
group quota rules, it will be the UNIX group name. For default rules (tree
or user or group), this will display `"*"` . If the UNIX user identifier,
UNIX group identifier, or Windows security identifier no longer exists or
if the identifier-to-name conversion fails, the target appears in numeric
form.
* Quota target ID - This field gives the target of the quota rule in
numeric form. For tree quota rules, it will be the qtree ID of the qtree.
For group quota rules, it will be the UNIX group identifier. For UNIX user
quota rules, it will be the UNIX user identifier. For Windows user quota
rules, it will be the Windows security identifier in its native format.
For default rules (tree or user or group), `"*"` will be displayed.
* Disk space used - For a default quota, the value is 0.
* Disk space limit
* Soft disk space limit
* Threshold for disk space limit
* Files used - For a default quota, the value is 0.
* File limit
* Soft file limit
* Quota specifier - For an explicit quota, this field shows how the quota
target was configured by the administrator using the volume quota policy
rule command. For a default quota, the field shows `"*"` . For a derived
tree quota, this field shows the qtree path. For a derived user and group
quota, the field is either blank or `"*"` .



The following parameters: `-soft` , `-soft-limit-thresholds` , `-target-
id` , `-thresholds` , `-fields` and `-instance` display different set of
fields listed above. For example, `-soft` will display the soft disk space
limit and soft file limit apart from other information. Similarly `-
```

```
target-id` will display the target in the numeric form.

A quota report is a resource intensive operation. If you run it on many
volumes in the cluster, it might take a long time to complete. A more
efficient way would be to view the quota report for a particular volume in
a Vserver.

Depending upon the quota rules configured for a volume, the quota report
for a single volume can be large. If you want to monitor the quota report
entry for a particular tree/user/group repeatedly, find the index of that
quota report entry and use the `-index` field to view only that quota
report entry. See the examples section for an illustration.

== 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.

| ``[-soft  ]``::
If this parameter is specified, the command display will include the soft
disk space limit and the soft file limit.

| ``[-soft-limit-thresholds  ]``::
If this parameter is specified, the command display will include the soft
disk space limit, threshold for disk space limit and soft file limit.

| ``[-target-id  ]``::
If this parameter is specified, the command will display the target of a
user or group quota rule in numeric form.

| ``[-thresholds  ]``::
If this parameter is specified, the command display will include the
threshold for disk space limit.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name::
If this parameter is specified, the command displays the quota report for
volumes in the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter is specified, the command displays the quota report for
the specified volume.
```

```
``[-index <integer>]`` - Index::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified index.

``[-tree <qtree name>]`` - Qtree Name::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified qtree name.

``[-quota-type {tree|user|group}]`` - Quota Type::
If this parameter is specified, the command displays the quota report for
the quota rules of the given type.

``[-quota-target <text>,...]`` - Quota Target::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified quota target.

``[-quota-target-id <text>,...]`` - Quota Target ID::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified quota target identifier.

``[-disk-used {<integer>[KB|MB|GB|TB|PB]}]`` - Disk Space Used::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified disk space used value.

``[-disk-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Disk Space Limit::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified disk space limit.

``[-files-used <integer>]`` - Files Used::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified files used value.

``[-file-limit <integer>]`` - Files Limit::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified file limit.

``[-threshold {<integer>[KB|MB|GB|TB|PB]}]`` - Disk Space Threshold::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified threshold for disk space limit.

``[-soft-disk-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Soft Disk Space
Limit::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified soft disk space limit.

``[-soft-file-limit <integer>]`` - Soft Files Limit::
```

```
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified soft file limit.

``[-quota-specifier <text>]`` - Quota Specifier::
If this parameter is specified, the command displays the quota report for
the quota rules that have the specified quota specifier.

``[-path <text>]`` - Path::
If this parameter is specified, the command will display the quota report
for the quota rules that are applicable for the file in the specified
path. The format of the path to the file should begin with /vol/<volume
name>/. The quota rules that are applicable typically consists of the tree
quota rule corresponding to the qtree in which the file resides within the
volume, user quota rule at the volume and qtree level corresponding to the
UNIX user identifier or the Windows security identifier associated with
the file and the group quota rule at the volume and qtree level
corresponding to the UNIX group identifier associated with the file.

``[-disk-used-pct-soft-disk-limit <percent_no_limit>]`` - Disk Space Used
% Soft Disk Space Limit::
If this parameter is specified, the command displays the quota report for
entries that have the specified percent utilization. The attribute value
is computed from `disk-used` and `soft-disk-limit` .

``[-disk-used-pct-threshold <percent_no_limit>]`` - Disk Space Used % Disk
Space Threshold::
If this parameter is specified, the command displays the quota report for
entries that have the specified percent utilization. The attribute value
is computed from `disk-used` and `threshold` .

``[-disk-used-pct-disk-limit <percent_no_limit>]`` - Disk Space Used %
Disk Space Limit::
If this parameter is specified, the command displays the quota report for
entries that have the specified percent utilization. The attribute value
is computed from `disk-used` and `disk-limit` .

``[-files-used-pct-soft-file-limit <percent_no_limit>]`` - Files Used %
Soft File Limit::
If this parameter is specified, the command displays the quota report for
entries that have the specified percent utilization. The attribute value
is computed from `files-used` and `soft-file-limit` .

``[-files-used-pct-file-limit <percent_no_limit>]`` - Files Used % File
Limit::
If this parameter is specified, the command displays the quota report for
entries that have the specified percent utilization. The attribute value
```

```
is computed from `files-used` and `file-limit` .

== Examples

The following example displays the quota report for all the volumes.
```

cluster1::> volume quota report
Vserver: vs0
----Disk---- ----Files----- Quota
Volume Tree Type ID Used Limit Used Limit Specifier
------- -------- ------ ------- ----- ----- ------ ------ ---------
vol2 tree * 0.00B 100MB 0 10000 *
vol2 vxw02 tree 3 0.00B 200MB 1 20000 vxw02
vol2 user * 0.00B 50MB 0 - *
vol2 vxw02 user sam,Engr\Sammy
0.00B 100MB 0 - sam
vol2 group * 0.00B 500MB 0 - *
vol2 q1 tree 1 1MB 100MB 2 10000 q1
vol2 q1 user * 0.00B 50MB 0 -
vol2 q1 group * 0.00B 500MB 0 -
vol2 q1 group root 1MB - 2 -
vol2 vxw01 tree 2 0.00B 100MB 1 10000 vxw01
vol2 vxw01 user * 0.00B 50MB 0 -
vol2 vxw01 group * 0.00B 500MB 0 -
vol2 vxw01 group root 0.00B - 1 -
vol2 vxw02 user * 0.00B 50MB 0 -
vol2 vxw02 group * 0.00B 500MB 0 -
vol2 vxw02 group root 0.00B - 1 -
vol2 vxw03 tree 4 0.00B 100MB 1 10000 vxw03
vol2 vxw03 user * 0.00B 50MB 0 -
vol2 vxw03 group * 0.00B 500MB 0 -
vol2 vxw03 group root 0.00B - 1 -
vol2 group root 1MB - 6 -
vol2 q1 user root,Engr\root
0.00B - 1 -
vol2 vxw01 user root,Engr\root
0.00B - 1 -
vol2 vxw02 user root,Engr\root
0.00B - 1 -
vol2 vxw03 user root,Engr\root
0.00B - 1 -
vol2 user root,Engr\root
0.00B - 5 -
vol2 user john,Engr\John
1MB 50MB 1 - *
vol2 q1 user john,Engr\John
1MB 50MB 1 -
28 entries were displayed.

```
The following example displays the quota report for the quota rules that
are applicable for the given path to a file.
```

```
cluster1::> volume quota report -path /vol/vol2/q1/file1
Vserver: vs0
----Disk---- ----Files----- Quota
Volume Tree Type ID Used Limit Used Limit Specifier
------- -------- ------ ------- ----- ----- ------ ------ ---------
vol2 q1 tree 1 1MB 100MB 2 10000 q1
vol2 q1 group root 1MB - 2 -
vol2 group root 1MB - 6 -
vol2 user john,Engr\John
1MB 50MB 1 - *
vol2 q1 user john,Engr\John
1MB 50MB 1 -
5 entries were displayed.
```

```
The following example displays the quota report with the target in the
numeric form for the given path to a file.
```

```
cluster1::> volume quota report -path /vol/vol2/q1/file1 -target-id
Vserver: vs0
----Disk---- ----Files----- Quota
Volume Tree Type ID Used Limit Used Limit Specifier
------- -------- ------ ------- ----- ----- ------ ------ ---------
vol2 q1 tree 1 1MB 100MB 2 10000 q1
vol2 q1 group 0 1MB - 2 -
vol2 group 0 1MB - 6 -
vol2 user 8017,S-1-5-21-3567637-1906459281-1427260136-60871
1MB 50MB 1 - *
vol2 q1 user 8017,S-1-5-21-3567637-1906459281-1427260136-60871
1MB 50MB 1 -
5 entries were displayed.
```

```
The following example shows how to monitor the quota report for a
particular user/tree/group. First, the quota report command is issued with
`-instance` to see the index field for the quota report entry we are
interested in. Next, the quota report is issued with the `-index` field
specified to fetch only that particular quota report entry repeatedly to
view the usage over time.
```

```
cluster1::> volume quota report -vserver vs0 -volume vol1 -quota-type user -quota-target john -tree q1
-instance
Vserver Name: vs0
Volume Name: vol1
Index: 10
Qtree Name: q1
Quota Type: user
Quota Target: john
Quota Target ID: 5433
Disk Space Used: 50.5MB
Disk Space Limit: 100MB
Files Used: 205
```

```
Files Limit: -
Disk Space Threshold: 95MB
Soft Disk Space Limit: 80MB
Soft Files Limit: -
Quota Specifier: john
Disk Space Used % Soft Disk Space Limit: 63%
Disk Space Used % Disk Space Threshold: 53%
Disk Space Used % Disk Space Limit: 51%
Files Used % Soft File Limit: -
Files Used % File Limit: -

cluster1::> volume quota report -vserver vs0 -volume vol1 -index 10
Vserver Name: vs0
Volume Name: vol1
Index: 10
Qtree Name: q1
Quota Type: user
Quota Target: john
Quota Target ID: 5433
Disk Space Used: 55MB
Disk Space Limit: 100MB
Files Used: 410
Files Limit: -
Disk Space Threshold: 95MB
Soft Disk Space Limit: 80MB
Soft Files Limit: -
Quota Specifier: john
Disk Space Used % Soft Disk Space Limit: 69%
Disk Space Used % Disk Space Threshold: 58%
Disk Space Used % Disk Space Limit: 55%
Files Used % Soft File Limit: -
Files Used % File Limit: -

cluster1::> volume quota report -vserver vs0 -volume vol1 -index 10
Vserver Name: vs0
Volume Name: vol1
Index: 10
Qtree Name: q1
Quota Type: user
Quota Target: john
Quota Target ID: 5433
Disk Space Used: 60.7MB
Disk Space Limit: 100MB
Files Used: 500
Files Limit: -
Disk Space Threshold: 95MB
Soft Disk Space Limit: 80MB
Soft Files Limit: -
Quota Specifier: john
Disk Space Used % Soft Disk Space Limit: 76%
Disk Space Used % Disk Space Threshold: 64%
Disk Space Used % Disk Space Limit: 61%
Files Used % Soft File Limit: -
Files Used % File Limit: -
```

```
[[ID1341c45c23870d12e1109f6a6c773835]]
= volume quota resize
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resize quotas for volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command resizes the quota limits for the quotas currently in effect
for the specified volume. It creates a job to resize quotas. You can
monitor the progress of the job by using the xref:{relative_path}job-
show.html[job show] and xref:{relative_path}job-watch-progress.html[job
watch-progress] commands.

NOTE: Quotas must be activated before quota limits can be resized.



== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the name of the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the name of the volume on which you are resizing the quota
limits and threshold.

``[-foreground <true>]`` - Foreground Process::
This optionally specifies whether the job created for resizing quotas runs
as a foreground process. The default setting is `false` (that is, the
operation runs in the background). When set to `true` , the command will
not return until the job completes.

== Examples

The following example resizes quotas on the volume named `vol1` , which
exists on Vserver `vs0` .
```

cluster1::> volume quota resize -vserver vs0 -volume vol1
[Job 34] Job is queued: Quota RESIZE Operation on vserver vs0 volume vol1.

```
The following example uses a 7G-compatible command to resize quotas on the
volume named `vol1` which exists on Vserver `vs0` .
```

cluster1::> vserver context vs0
vs0::> quota resize vol1
[Job 35] Job is queued: Quota RESIZE Operation on vserver vs0 volume vol1.

```
== Related Links

* xref:{relative_path}job-show.html[job show]
* xref:{relative_path}job-watch-progress.html[job watch-progress]




[[ID78cd0c31034c5a23f4ca9e77a1fe9552]]
= volume quota show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display quota state for volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays information about quotas for online volumes. The
command output depends on the parameters specified with the command.
Quotas can only be administered on FlexVol volumes. If no parameters are
specified, the command displays the following information for all volumes:

* Vserver name
* Volume name
* Quota state - quota state for this volume. The possible values are as
follows:

  * `off` - this state indicates that quotas are deactivated for the
volume.
  * `on` - this state indicates that quotas are activated for the volume.
```

```
  * `initializing` - this state indicates that quotas are being
initialized for the volume.
  * `resizing` - this state indicates that quota limits are being resized
for the volume.
  * `corrupt` - this state indicates that quotas are corrupt for this
volume.
  * `mixed` - this state may only occur for FlexGroup volumes, and
indicates that the constituent volumes are not all in the same state.


* Scan status - percentage of the files in the volume scanned by the quota
scanner that runs as part of activating quotas.
* Last error - most recently generated error message as part of the last
quota operation (`on` or `resize` ). Present only if an error has been
generated.



To display detailed information about all volumes, run the command with
the -instance parameter. The detailed view provides all of the information
in the previous list and the following additional information:

* Logging messages - whether quota exceeded syslog/EMS messages are logged
or not. For volumes where the quota logging parameter is set to `on` ,
quota exceeded messages are generated when a NFS/CIFS operation or any
internal ONTAP operation is being prevented because the quota disk usage
is exceeding the disk limit or the quota file usage is exceeding the file
limit. For quotas where the logging parameter is set to `off` , no quota
exceeded messages are generated.
* Logging interval - frequency with which quota exceeded messages are
logged. This parameter only applies to volumes that have the logging
parameter set to `on` .
* Sub status - additional status about quotas for this volume. Following
are the possible values reported:

  * `upgrading` - this indicates that the quota metadata format is being
upgraded from an older version to a newer version for the volume.
  * `setup` - this indicates that the quotas are being setup on the
volume.
  * `transferring rules` - this indicates that the quota rules are being
transferred to the volume.
  * `scanning` - this indicates that the quota filesystem scanner is
currently running on the volume.
  * `finishing` - this indicates that the quota `on` or `resize` operation
is in the final stage of the operation.
  * `done` - this indicates that the quota operation is finished.
```

```
  * `none` - this indicates that there is no additional status.


* All errors - collection of all the error messages generated as part of
the last quota operation (`on` or `resize` ) since the volume was online.



* User quota enforced (advanced privilege only) - indicates whether there
are user quota rules being enforced.
* Group quota enforced (advanced privilege only)- indicates whether there
are group quota rules being enforced.
* Tree quota enforced (advanced privilege only) - indicates whether there
are tree quota rules being enforced.



== 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.

| ``[-logmsg  ]``::
If this parameter is specified, the command displays whether quota
exceeded messages are logged and the logging interval for the quota
messages.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name::
If this parameter is specified, the command displays information for the
volumes in the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter is specified, the command displays information for the
specified volume.

``[-state <quota_state>]`` - Quota State::
If this parameter is specified, the command displays information for the
volumes that have the specified quota state.

``[-scan-status <percent>]`` - Scan Status::
If this parameter is specified, the command displays information about the
```

```
volumes whose scan-status matches the specified percentage value. The scan
status is displayed in the format [0-100]%.

``[-logging {on|off}]`` - Logging Messages::
If this parameter is specified, the command displays information about the
volumes that have the specified logging setting.

``[-logging-interval <text>]`` - Logging Interval::
If this parameter is specified, the command displays information about the
volumes that have the specified quota logging interval.

``[-sub-status <text>]`` - Sub Quota Status::
If this parameter is specified, the command displays information about the
volumes that have the specified quota sub-status.

``[-last-error <text>]`` - Last Quota Error Message::
If this parameter is specified, the command displays information about the
volumes whose last error matches the specified error message.

``[-errors <text>]`` - Collection of Quota Errors::
If this parameter is specified, the command displays information about the
volumes whose collection of errors match the specified error message.

``[-is-user-quota-enforced {true|false}]`` - User Quota enforced
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information about the
volumes that have the specified value for this field.

``[-is-group-quota-enforced {true|false}]`` - Group Quota enforced
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information about the
volumes that have the specified value for this field.

``[-is-tree-quota-enforced {true|false}]`` - Tree Quota enforced
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information about the
volumes that have the specified value for this field.

== Examples

The following example displays information about all volumes on the
Vserver named `vs0` .
```

cluster1::> volume quota show -vserver vs0
Scan
Vserver Volume State Status

--------- ---------- --------- ------
vs0 root_vs0 off -
vs0 vol1 off -
Last Error: Volume vol1 has no valid quota rules
vs0 vol2 on -
vs0 vol3 initializing 30%
4 entries were displayed.

```
The following example displays the logging information for all the
volumes.
```

cluster1::> volume quota show -logmsg
Logging Logging
Vserver Volume State Message Interval
--------- ---------- --------- ------- --------
vs0 root_vs0 off - -
vs0 vol1 off - -
vs0 vol2 on on 1h
vs0 vol3 on on 1h
4 entries were displayed.

```
The following example displays detailed information in advanced privilege
for a volume `vol1` , which exists on Vserver `vs0`
```

cluster1::> set advanced

Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y

cluster1::*> volume quota show -instance -vserver vs0 -volume vol1
Vserver Name: vs0
Volume Name: vol1
Quota State: on
Scan Status: -
Logging Messages: on
Logging Interval: 1h
Sub Quota Status: none
Last Quota Error Message: -
Collection of Quota Errors: -
User Quota enforced: true
Group Quota enforced: false
Tree Quota enforced: false

```
The following example displays the detailed information when quotas are
upgrading for volume `vol1` , which exists on Vserver `vs0` .
```

cluster1::> volume quota show -instance -vserver vs0 -volume vol1

Vserver Name: vs0
Volume Name: vol1
Quota State: initializing
Scan Status: 3%
Logging Messages: -
Logging Interval: -
Sub Quota Status: upgrading
Last Quota Error Message: -
Collection of Quota Errors: -

```
The following example displays the "Last Quota Error Message" and the
"Collection of Quota Errors" for volume `vol1` , which exists on Vserver
`vs0`
```

cluster1::> volume quota show -instance -vserver vs0 -volume vol1
Vserver Name: vs0
Volume Name: vol1
Quota State: on
Scan Status: -
Logging Messages: on
Logging Interval: 1h
Sub Quota Status: none
Last Quota Error Message: second definition for user2 (type:user target:user2,user4 qtree:"").
Collection of Quota Errors: second definition for user1 (type:user target:user1,user3 qtree:"").
second definition for user2 (type:user target:user2,user4 qtree:"").

```
[[IDac4725bd3e4dafc5c4a65ded13d3914a]]
= volume quota policy copy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Copy a quota policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command copies a quota policy and the rules it contains.You must
enter the following information to copy a quota policy:

* Vserver name
* Policy name
* New policy name



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver from which you are copying the quota
policy.

``-policy-name <text>`` - Policy Name::
This parameter specifies the name of the quota policy you are copying.

``-new-policy-name <text>`` - New Policy Name::
This parameter specifies the name of the new quota policy you are copying
to. The new name cannot have more than 32 characters.

== Examples

The following example copies a quota policy named quota_policy_0 on
Vserver vs0. It is copied to quota_policy_1.
```

cluster1::> volume quota policy copy -vserver vs0 -policy-name quota_policy_0 -new-policy-name quota_policy_1

```
[[ID72249169fecde50ef5f5348e49dfea4e]]
= volume quota policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a quota policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

A quota policy is collection of quota rules for all the volumes in a
specific Vserver. This command creates a quota policy for a specific
Vserver. Multiple quota policies can be created for a Vserver, but only
one of them can be assigned to the Vserver. A Vserver can have a maximum
of five quota policies. If five quota policies already exist, the command
fails and a quota policy must be deleted before another quota policy can
be created.

When you turn on quotas for a volume, the quota rules to be enforced on
that volume will be picked from the quota policy that is assigned to the
Vserver containing that volume. The quota policy for clustered volumes is
equivalent to the /etc/quotas file in 7G.

You must enter the following information to create a quota policy:

* Vserver name
* Policy name



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver for which you are creating the quota
policy. You can create a quota policy only for a data Vserver. Quota
policies cannot be created for a node or admin Vserver.

``-policy-name <text>`` - Policy Name::
This parameter specifies the name of the quota policy you are creating.
The quota policy name cannot be more than 32 characters long and must be
unique within the Vserver.
```

```
== Examples

The following example creates a quota policy named quota_policy_0 on
Vserver vs0.
```

cluster1::> volume quota policy create -vserver vs0 -policy-name quota_policy_0

```
[[ID696abfca729cdccd6bbcbb8d17ffef43]]
= volume quota policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a quota policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command deletes a quota policy and all the rules it contains. The
policy can be deleted only when it is not assigned to the Vserver.You must
enter the following information to delete a quota policy:

* Vserver name
* Policy name



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver containing the quota policy you want
to delete.

``-policy-name <text>`` - Policy Name::
This parameter specifies the name of the quota policy you want to delete.

== Examples

The following example deletes a quota policy named quota_policy_5 on
Vserver vs0.
```

cluster1::> volume quota policy delete -vserver vs0 -policy-name quota_policy_5

```
[[IDc77467c4c23893cd646847625854978c]]
= volume quota policy rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a quota policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command renames a quota policy. You must enter the following
information to rename a quota policy:

* Vserver name
* Policy name
* New policy name



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver containing the quota policy you want
to rename.

``-policy-name <text>`` - Policy Name::
This parameter specifies the name of the quota policy you are renaming.

``-new-policy-name <text>`` - New Policy Name::
This parameter specifies the new name of the quota policy. The new name
cannot be more than 32 characters long.

== Examples

The following example renames a quota policy named quota_policy_0 on
Vserver vs0. The policy's new name is quota_policy_1.
```

cluster1::> volume quota policy rename -vserver vs0 -policy-name quota_policy_0 -new-policy-name quota_policy_1

```
[[IDec32113be4af816f5ceb88b24eb08e42]]
= volume quota policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the quota policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays information about quota policies. The command
displays the following information about all quota policies:

* Vserver name
* Policy name
* When the policy was last modified



== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays information about the
quota policies for the specified Vserver.

``[-policy-name <text>]`` - Policy Name::
If this optional parameter is specified, the command displays information
about quota policies that match the specified name.

``[-last-modified <MM/DD/YYYY HH:MM:SS>]`` - Last Modified::
If this optional parameter is specified, the command displays information
about quota policies with the last modified time that match the given
```

```
time.

== Examples

The following example displays information about all quota policies.
```

cluster1::> volume quota policy show
Vserver Policy Name Last Modified
--- --- ---
vs0 quota_policy_vs0 10/16/2008 17:40:05
vs1 quota_policy_vs1 10/16/2008 17:47:45
vs2 quota_policy_vs2 10/16/2008 17:44:13
vs3 quota_policy_vs3 10/16/2008 17:44:13
4 entries were displayed.

```
[[ID335f1372659361ded4227b13e2bf9d50]]
= volume quota policy rule create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new quota rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates a quota policy rule. You must enter the following
information to create a quota policy rule:

* Vserver name
* Quota policy name
* Volume name
* Quota target type
* Target to which the rule applies
* Qtree to which the rule applies



You can optionally specify the following additional attributes for the
quota policy rule:

* User mapping
* Hard disk limit
```

```
* Hard file limit
* Threshold for disk limit
* Soft disk limit
* Soft file limit



NOTE: For a new quota policy rule to get enforced on the volume, you
should create the rule in the quota policy assigned to the Vserver.
Additionally, a quota off and on or a quota resize operation must be done
using the "link:volume-quota-modify.html[volume quota modify] " command.



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver containing the quota policy for which
you are creating a rule.

``-policy-name <text>`` - Policy Name::
This parameter specifies the name of the quota policy in which you are
creating a rule.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume for which you are creating
a rule.

``-type {tree|user|group}`` - Type::
This parameter specifies the quota target type of the rule you are
creating.

``-target <text>`` - Target::
This parameter specifies the target to which the quota policy rule
applies. For default quota rules, this parameter should be specified as
"". For explicit tree quotas rules, this parameter should indicate the
qtree name. For explicit user quota rules, this parameter can contain UNIX
user name, UNIX user identifier, Windows user name, Windows Security
Identifier or a path to an existing object within the volume. If a name
contains a space, enclose the entire value in quotes. A UNIX user name
cannot include a backslash (\) or an @ sign; user names with these
characters are treated as Windows names. For multi-user quotas, this
parameter can contain multiple user targets separated by a comma. For
explicit group quota rules, this parameter can contain UNIX group name or
UNIX group identifier or a path to an existing object within the volume.
When a path is specified as the target, it should be of the format
/vol/<vol-name>/<path to file from volume root> where the volume matches
```

```
that of the `-volume` parameter.

``-qtree <qtree name>`` - Qtree Name::
This parameter specifies the name of the qtree to which the quota rule
applies. This parameter is not applicable for tree type rules. For user or
group type rules at the volume level, this parameter should contain "".

``[-user-mapping {on|off}]`` - User Mapping::
This parameter optionally specifies if user mapping needs to be performed
for a user quota rule. If this parameter is "on", the UNIX user name
specified as the quota target will be mapped to the corresponding Windows
user name or vice-versa and quota accounting will be performed for the
users together. The mapping will be obtained as configured in "`vserver
name-mapping` ".
+
Note that this parameter can be specified only for quota policy rules of
type user. A value of "on" can be specified for this parameter only if the
quota target is a UNIX user name or a Windows user name and cannot be
specified for multi-user quota targets.

``[-disk-limit {<size>|-}]`` - Disk Limit::
This parameter optionally specifies a hard limit for the disk space that
can be consumed by the quota target. The default unit for the disk limit
is assumed to be Kilobytes if no units are specified. When the hard disk
space limit is reached, no additional disk space can be consumed by the
specified target. The value that you specify for this parameter should be
greater than or equal to the threshold and soft disk limit. A disk limit
of unlimited can be specified with a "-" for this parameter or by not
specifying this parameter and will be indicated by a "-". The maximum
value is 1,125,899,906,842,620 KB, which is approximately 1,023 PB. The
value should be a multiple of 4 KB. If it is not, this field can appear
incorrect in quota reports. This happens because the field is always
rounded up to the nearest multiple of 4 KB to match disk space limits,
which are translated into 4-KB disk blocks. The value can be larger than
the amount of disk space available in the volume.

``[-file-limit {<integer>|-}]`` - Files Limit::
This parameter optionally specifies a hard limit for the number of files
permitted on the quota target. When the hard number of files limit is
reached, no additional files can be created by the specified target. The
value that you specify for this parameter should be greater than or equal
to the soft file limit. A file limit of unlimited can be specified with a
"-" for this parameter or by not specifying this parameter and will be
indicated by a "-". The maximum value is 4,294,967,295.

``[-threshold {<size>|-}]`` - Threshold for Disk Limit::
```

This parameter optionally specifies the disk limit threshold for the quota target. The default unit for the disk limit threshold is assumed to be Kilobytes if no units are specified. When the disk limit threshold is exceeded, a console message, EMS events, and SNMP traps are generated. The value that you specify for this parameter should be greater than or equal to the soft disk limit and equal to or less than the disk limit. A threshold of unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 1,125,899,906,842,620 KB, which is approximately 1,023 PB. The value should be a multiple of 4 KB. If it is not, this field can appear incorrect in quota reports. This happens because the field is always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks.

``[-soft-disk-limit {<size>|-}]`` - Soft Disk Limit::
This parameter optionally specifies a soft limit for the disk space that can be consumed by the quota target. The soft disk limit indicates that the hard limit for the disk space will soon be exceeded. The default unit for the soft disk limit is assumed to be Kilobytes if no units are specified. When the soft limit for the disk space is exceeded, a console message, EMS events and SNMP traps are generated. The same happens when the disk space used goes below the specified limit. The value that you specify for this parameter should be equal to or less than the threshold and the disk limit. A soft disk limit of unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 1,125,899,906,842,620 KB, which is approximately 1,023 PB. The value should be a multiple of 4 KB. If it is not, this field can appear incorrect in quota reports. This happens because the field is always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks.

``[-soft-file-limit {<integer>|-}]`` - Soft Files Limit::
This parameter optionally specifies a soft limit for the number of files permitted on the quota target. The soft file limit indicates that the hard limit for the number of files will soon be exceeded. When the soft limit for the number of files is exceeded, a console message, EMS events and SNMP traps are generated. The same happens when the files used goes below the specified limit. The value that you specify for this parameter should be equal to or less than the file limit. A soft file limit of unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 4,294,967,295.

== Examples

The following example creates a default tree quota rule for volume

```
``_vol0_`` in Vserver ``_vs0_`` and in the quota policy named
``_quota_policy_0_`` . This quota policy applies to all qtrees on volume
``_vol0_`` .
```

cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type tree -target ""

```
The following example creates a quota policy rule for volume ``_vol0_`` in
Vserver ``_vs0_`` and in the quota policy named ``_quota_policy_0_`` .
This quota policy applies to the UNIX user ``_myuser_`` for a qtree named
``_qtree1_`` on volume ``_vol0_`` with a disk limit of 20 Gigabytes, soft
disk limit of 15.4 Gigabytes and threshold limit of 15.4 Gigabytes. User
mapping is turned on for this rule.
```

cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target myuser
-qtree qtree1 -user-mapping on -disk-limit 20GB -soft-disk-limit 15.4GB
-threshold 15.4GB

```
The following example creates a quota policy rule for volume ``_vol0_`` in
Vserver ``_vs0_`` and in the quota policy named ``_quota_policy_0_`` .
This quota policy applies to the Windows user ``_DOMXYZ\myuser_`` for a
qtree named ``_qtree1_`` on volume ``_vol0_`` with a file limit of
``_40000_`` and a soft file limit of ``_26500_`` . User mapping is turned
on for this rule.
```

cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target DOMXYZ\myuser
-qtree qtree1 -user-mapping on -file-limit 40000 -soft-file-limit 26500

```
The following example creates a quota policy rule for volume ``_vol0_`` in
Vserver ``_vs0_`` and in the quota policy named ``_quota_policy_0_`` .
This quota policy applies to the UNIX user identifier ``_12345_`` for a
qtree named ``_qtree1_`` on volume ``_vol0_`` .
```

cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target 12345
-qtree qtree1

```
The following example creates a quota policy rule for volume ``_vol0_`` in
Vserver ``_vs0_`` and in the quota policy named ``_quota_policy_0_`` .
This quota policy applies to the Windows Security Identifier ``_S-123-456-
789_`` for a qtree named ``_qtree1_`` on volume ``_vol0_`` .
```

cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user
-target S-123-456-789 -qtree qtree1

```
The following example creates a quota policy rule for volume ``_vol0_`` in
Vserver ``_vs0_`` and in the quota policy named ``_quota_policy_0_`` .
This quota policy applies to the UNIX group ``_engr_`` for a qtree named
``_qtree1_`` on volume ``_vol0_`` .
```

cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type group -target engr
-qtree qtree1

```
The following example creates a quota policy rule for volume ``_vol0_`` in
Vserver ``_vs0_`` and in the quota policy named ``_quota_policy_0_`` .
This quota policy applies to the user who is the owner of the file
``_/vol/vol0/qtree1/file1.txt_`` for qtree ``_qtree1_`` on volume
``_vol0_`` .
```

cluster1::> volume quota policy rule create -vserver vs0 -policy-name
quota_policy_0 -volume vol0 -type user -target /vol/vol0/qtree1/file1.txt
-qtree qtree1

```
The following example creates a quota policy rule for volume ``_vol0_`` in
Vserver ``_vs0_`` and in the quota policy named ``_quota_policy_0_`` .
This quota policy applies to the users specified in the target for qtree
``_qtree1_`` on volume ``_vol0_`` .
```

cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user
-target user1,DOMXYZ\user2,23457,S-126-098-567,/vol/vol0/qtree1/file2.txt
-qtree qtree1

```
== Related Links

* xref:{relative_path}volume-quota-modify.html[volume quota modify]




[[ID1141658da7299834f26820f26ba0a20c]]
= volume quota policy rule delete
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing quota rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume quota policy rule delete` command deletes a quota policy rule.
You must enter the following information to delete a quota policy rule:

* Vserver name
* Quota policy name
* Volume name
* Quota target type
* Target to which the rule applies
* Qtree to which the rule applies



NOTE: If the rule being deleted belongs to the quota policy that is
currently assigned to the Vserver, enforcement of the rule on the volume
must be terminated by performing a quota off and on or a quota resize
operation using the "link:volume-quota-modify.html[volume quota modify] "
command.



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver containing the quota policy for which
you are deleting a rule.

``-policy-name <text>`` - Policy Name::
This parameter specifies the name of the quota policy in which you are
deleting a rule.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume for which you are deleting
a rule.

``-type {tree|user|group}`` - Type::
This parameter specifies the quota target type for the rule.
```

```
``-target <text>`` - Target::
This parameter specifies the target to which the quota policy rule
applies.

``-qtree <qtree name>`` - Qtree Name::
This parameter specifies the name of the qtree for which you are deleting
a rule.

== Examples

The following example deletes a quota policy rule on Vserver vs1 for the
quota policy named quota_policy_1. This quota policy applies to the group
named engr for the qtree named qtree1 on volume vol1.
```

cluster1::> volume quota policy rule delete -vserver vs1
-policy-name quota_policy_1 -volume vol1 -type group -target engr
-qtree qtree1

```
== Related Links

* xref:{relative_path}volume-quota-modify.html[volume quota modify]




[[ID d7d1ac786787a4564e0caba071bd9ddf]]
= volume quota policy rule modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an existing quota rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command can be used to modify the following attributes of a quota
policy rule:

* User mapping
* Hard disk limit
```

```
* Hard file limit
* Threshold for disk limit
* Soft disk limit
* Soft file limit



NOTE: If the rule being modified belongs to the quota policy that is
currently assigned to the Vserver, rule enforcement on the volume must be
enabled by performing a quota off and on or a quota resize operation using
the "link:volume-quota-modify.html[volume quota modify] " command.



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver containing the quota policy for which
you are modifying a rule.

``-policy-name <text>`` - Policy Name::
This parameter specifies the name of the quota policy in which you are
modifying a rule.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume for which you are
modifying a rule.

``-type {tree|user|group}`` - Type::
This parameter specifies the quota target type for the rule you are
modifying.

``-target <text>`` - Target::
This parameter specifies the target to which the quota policy rule
applies. If the target is a user, the user ID or username must be the same
one that was used to create the quota. The same restriction is there for
both group ID or groupname and Windows SID or Windows account name.

``-qtree <qtree name>`` - Qtree Name::
This parameter specifies the name of the qtree to which the quota policy
rule applies.

``[-user-mapping {on|off}]`` - User Mapping::
This parameter optionally modifies the user mapping for a user quota rule.
The value for this parameter can be modified only for quota policy rules
of type user. A value of "on" can be specified for this parameter only if
the quota target is a unix user name or a Windows user name and cannot be
```

```
specified for multi-user quota targets. If this parameter is "on", the
unix user name specified as the quota target will be mapped to the
corresponding Windows user name or vice-versa and quota accounting will be
performed for the users together.

``[-disk-limit {<size>|-}]`` - Disk Limit::
This parameter optionally modifies the hard limit for the disk space that
can consumed by the quota target. The default unit for the disk limit is
assumed to be Kilobytes if no units are specified. The value that you
specify for this parameter should be greater than or equal to the
threshold and soft disk limit. A disk limit of unlimited can be specified
with a "-" for this parameter.

``[-file-limit {<integer>|-}]`` - Files Limit::
This parameter optionally modifies the hard limit for the number of files
permitted on the quota target. The value that you specify for this
parameter should be greater than or equal to the soft file limit. A file
limit of unlimited can be specified with a "-" for this parameter

``[-threshold {<size>|-}]`` - Threshold for Disk Limit::
This parameter optionally modifies the disk limit threshold for the quota
target. The default unit for the disk limit threshold is assumed to be
Kilobytes if no units are specified. The value that you specify for this
parameter should be greater than or equal to the soft disk limit and equal
to or less than the disk limit. A threshold limit of unlimited can be
specified with a "-" for this parameter.

``[-soft-disk-limit {<size>|-}]`` - Soft Disk Limit::
This parameter optionally modifies the soft limit for the disk space that
can be consumed by the quota target. The default unit for the soft disk
limit is assumed to be Kilobytes if no units are specified. The value that
you specify for this parameter should be equal to or less than the
threshold and the disk limit. A soft disk limit of unlimited can be
specified with a "-" for this parameter.

``[-soft-file-limit {<integer>|-}]`` - Soft Files Limit::
This parameter optionally modifies the soft limit for the number of files
permitted on the quota target. The value that you specify for this
parameter should be equal to or less than the file limit. A soft file
limit of unlimited can be specified with a "-" for this parameter.

== Examples

The following example modifies a quota policy rule for the quota policy
named quota_policy_0. This quota policy exists on Vserver vs0 and applies
to the user named myuser for qtree named qtree1 on volume vol0. The user
```

```
mapping is turned on, the hard disk limit is set to 20 GB and the hard
file limit is set to 100,000 files.
```

cluster1::> volume quota policy rule modify -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target myuser
-qtree qtree1 -user-mapping on -disk-limit 20GB -file-limit 100000

```
== Related Links

* xref:{relative_path}volume-quota-modify.html[volume quota modify]




[[ID675d74fedbe83d7cc0f9c3019ad6f151]]
= volume quota policy rule show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the quota rules

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays the following information about quota policy rules
by default.

* Vserver name
* Quota policy name
* Volume name
* Type of quota policy rule
* Target of the quota policy rule
* Qtree name
* User mapping
* Hard disk limit
* Soft disk limit
* Hard file limit
* Soft file limit
* Threshold for disk limit
```

```
== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays information about
quota rules for the quotas contained on volumes on the specified Vserver.

``[-policy-name <text>]`` - Policy Name::
If this parameter is specified, the command displays information about
quota rules for the specified quota policy.

``[-volume <volume name>]`` - Volume Name::
If this parameter is specified, the command displays information about
quota rules for the quota policy associated with the specified volume.

``[-type {tree|user|group}]`` - Type::
If this parameter is specified, the command displays information about
quota rules for the specified quota type.

``[-target <text>]`` - Target::
If this parameter is specified, the command displays information about
quota rules for the specified target.

``[-qtree <qtree name>]`` - Qtree Name::
If this parameter is specified, the command displays information about
quota rules for the specified qtree.

``[-user-mapping {on|off}]`` - User Mapping::
If this parameter is specified, the command displays information about
quota rules having the specified user-mapping value.

``[-disk-limit {<size>|-}]`` - Disk Limit::
If this parameter is specified, the command displays information about
quota rules having the specified hard disk limit.

``[-file-limit {<integer>|-}]`` - Files Limit::
If this parameter is specified, the command displays information about
quota rules having the specified hard file limit.
```

```
``[-threshold {<size>|-}]`` - Threshold for Disk Limit::
If this parameter is specified, the command displays information about
quota rules having the specified disk limit threshold.

``[-soft-disk-limit {<size>|-}]`` - Soft Disk Limit::
If this parameter is specified, the command displays information about
quota rules having the specified soft disk limit.

``[-soft-file-limit {<integer>|-}]`` - Soft Files Limit::
If this parameter is specified, the command displays information about
quota rules having the specified soft file limit.

== Examples

The following example displays information about all the quota policy
rules in a cluster. There is one user rule that exists on Vserver vs0 for
the quota policy named quota_policy_0. This quota policy applies to the
user named myuser for qtree named qtree0 on volume vol0.
```

cluster1::> volume quota policy rule show
Vserver: vs0 Policy: quota_policy_0 Volume: vol0
Soft Soft
User Disk Disk Files Files
Type Target Qtree Mapping Limit Limit Limit Limit Threshold
----- ------ ------ ------- ------- ------- ------ ------- ---------
tree myuser qtree0 on 20GB 18GB 100000 80000 16GB

```
[[ID4453951476802357d92ae477f50ed7e4]]
= volume quota policy rule count show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display count of quota rules

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 This command displays various counts of quota policy rules defined within
a quota policy. By default, the subtotal for each volume is displayed.
Optionally, the command can provide the total rule count across the entire
```

```
quota policy or detailed subtotals organized by qtree and quota rule type.

== 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.

| ``[-detail  ]``::
Displays rule count subtotals for each quota rule type. The subtotals for
each type are computed for a specific volume and qtree.

| ``[-hierarchy  ]``::
Displays rule count subtotals in hierarchical format with subtotals at the
quota policy, volume, qtree, and quota rule type levels.

| ``[-total  ]``::
Displays the total rule count for each Vserver and quota policy.

| ``[-instance  ]`` }::
Displays detailed information about all fields.

``[-vserver <vserver name>]`` - Vserver::
Displays quota rule counts for the specified Vserver.

``[-policy-name <text>]`` - Policy Name::
Displays quota rule counts for the specified quota policy.

``[-volume <volume name>]`` - Volume Name::
Displays quota rule counts for the specified volume.

``[-qtree <qtree name>]`` - Qtree Name::
Displays quota rule counts for the specified qtree.

``[-type {tree|user|group}]`` - Type::
Displays quota rule counts for the specified quota rule type.

``[-count-where-policy-volume-qtree-type <integer>]`` - Qtree/Type
Subtotal::
Subtotal of rules matching the given Vserver, quota policy, volume, qtree,
and quota rule type. If specified as input, only matching totals are
displayed.

``[-count-where-policy-volume-qtree <integer>]`` - Qtree Subtotal::
Subtotal of rules matching the given Vserver, quota policy, volume, and
qtree. All quota rule types are included. If specified as input, only
```

```
matching totals are displayed.

``[-count-where-policy-volume-type <integer>]`` - Volume/Type Subtotal::
Subtotal of rules matching the given Vserver, quota policy, volume, and
quota rule type. All qtrees are included. If specified as input, only
matching totals are displayed.

``[-count-where-policy-volume <integer>]`` - Volume Subtotal::
Subtotal of rules matching the given Vserver, quota policy, and volume.
All qtrees and quota rule types are included. If specified as input, only
matching totals are displayed.

``[-count-where-policy-type <integer>]`` - Policy/Type Subtotal::
Subtotal of rules matching the given Vserver, quota policy, and quota rule
type. All volumes and qtrees are included. If specified as input, only
matching totals are displayed.

``[-count-where-policy <integer>]`` - Policy Total::
Total rule count matching the given Vserver and quota policy. All volumes,
qtrees, and quota rule types are included. If specified as input, only
matching totals are displayed.

== Examples

The following example shows quota rule counts for Vserver ``_vs0_`` ,
quota policy ``_default_`` . The total number of rules in quota policy
``_default_`` is 7500. There are two volumes with quota rules configured.
Volume ``_volume0_`` has a total of 1000 rules, and ``_volume1_`` has a
total of 6500 rules.
```

cluster1::> volume quota policy rule count show -vserver vs0 -policy-name default

Vserver: vs0 Policy: default
Rule
Volume Count
--- ---
volume0 1000
volume1 6500
2 entries were displayed.

```
:leveloffset: -1


= volume reallocation commands

:leveloffset: +1
```

```
[[ID91e7b4ea2f57d6681cdee86b2a8ac6e1]]
= volume reallocation measure
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start reallocate measure job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Performs a measure-only reallocation check on a LUN, NVMe namespace, file,
or volume. At the end of each check, the system logs the optimization
results in the Event Message System (EMS). If you use the `logfile` , the
system records detailed information about the LUN, NVMe namespace, file,
or volume layout in the log file. To view previous measure-only
reallocation checks, use the xref:{relative_path}volume-reallocation-
show.html[volume reallocation show] command.

NOTE: This command is not supported for FlexGroup volumes or FlexGroup
constituents.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver.

``-path <text>`` - Path::
Specifies the path of the reallocation for a LUN, NVMe namespace, file, or
volume.

{ ``[-i, -interval <text>]`` - Interval Schedule::
Specifies the reallocation scan interval in
+
  * m for minutes
  * h for hours
  * d for days

+
For example, `30m` is a 30 minute interval. The countdown to the next scan
```

```
begins after the first scan is complete.
+
The default interval is 24 hours.

| ``[-o, -once <true>]`` - Once }::
Specifies that the job runs once and then is automatically removed from
the system when set to true. If you use this command without specifying
this parameter, its effective value is false and the reallocation scan
runs as scheduled. If you enter this parameter without a value, it is set
to true and a reallocation scan runs once.

``[-l, -logpath <text>]`` - Log Path::
Specifies the path for reallocation logs.

``[-t, -threshold <integer>]`` - Threshold::
Specifies the threshold when a LUN, NVMe namespace, file, or volume is
considered unoptimized and a reallocation should be performed. Once the
threshold is reached, the system creates a diagnostic message that
indicates that a reallocation might improve performance.
+
The threshold range is from 3 (the layout is moderately optimized) to 10
(the layout is not optimal). The threshold default is 4.

== Examples
```

cluster1::> volume reallocation measure -path /vol/vol2 -once
[Job 167] Job is queued: Reallocate Job.

```
Performs a one-time, measure-only reallocation scan on volume vol2.

== Related Links

* xref:{relative_path}volume-reallocation-show.html[volume reallocation
show]




[[IDaa6fa50153cd5a278da7ccf697c9f2a9]]
= volume reallocation off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable reallocate jobs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Disables all reallocation jobs globally in a cluster. After you use this
command, you cannot start or restart any reallocation jobs. All jobs that
are executing when you use this command are stopped. You must use the
`reallocate on` command to enable or restart reallocation jobs globally in
a cluster.

NOTE: This command is not supported for FlexGroup volumes or FlexGroup
constituents.



== Examples
```

cluster1::> volume reallocation off

```
Disables all reallocation jobs globally in a cluster.




[[ID2ed5bdba90915e59fb5600c3d6dbf5a1]]
= volume reallocation on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable reallocate jobs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Globally enables all reallocation jobs in a cluster. You must globally
enable reallocation scans in the cluster before you can run a scan or
schedule regular scans. Reallocation scans are disabled by default.

NOTE:

* This command is not supported for FlexGroup volumes or FlexGroup
constituents.
*





== Examples
```

cluster1::> volume reallocation on

```
Globally enables all reallocation jobs on a cluster.




[[ID2512270e4a4e6822a7aa098b54f46388]]
= volume reallocation quiesce
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Quiesce reallocate job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Temporarily stops any reallocation jobs that are in progress. When you use
this command, the persistent state is saved. You can use the
xref:{relative_path}volume-reallocation-restart.html[volume reallocation
restart] command to restart a job that is quiesced.

There is no limit to how long a job can remain in the quiesced state.

NOTE: This command is not supported for FlexGroup volumes or FlexGroup
constituents.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver.

``-path <text>`` - Path::
Specifies the file path of the LUN, NVMe namespace, file, or volume that
you want to stop temporarily.

== Examples
```

cluster1::> volume reallocation quiesce /vol/vol2
2 entries were acted on.

```
Temporarily stops all reallocation jobs on volume vol2.

== Related Links

* xref:{relative_path}volume-reallocation-restart.html[volume reallocation
restart]




[[ID07a593810b3246d12a2a0cc75f367782]]
= volume reallocation restart
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restart reallocate job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Starts a reallocation job. Use this command to start a quiesced
(temporarily stopped) job or a scheduled scan that is idle.

NOTE: This command is not supported for FlexGroup volumes or FlexGroup
constituents.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver.

``-path <text>`` - Path::
Specifies the file path of the LUN, NVMe namespace, file, or volume on
which you want to restart reallocation scans.

``[-i, -ignore-checkpoint <true>]`` - Ignore Checkpoint::
Restarts the job at the beginning when set to true. If you use this
command without specifying this parameter, its effective value is false
and the job starts the scan at the point where it stopped. If you specify
this parameter without a value, it is set to true and the scan restarts at
```

```
the beginning.

== Examples
```

cluster1::> volume reallocation restart /vol/vol2
2 entries were acted on.

```
Restarts two reallocation jobs on volume vol2.




[[ID82f7990f72d669e28aa345557bbd5e69]]
= volume reallocation schedule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify schedule of reallocate job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Schedules a reallocation scan for an existing reallocation job. If the
reallocation job does not exist, use the xref:{relative_path}volume-
reallocation-start.html[volume reallocation start] command to define a
reallocation job.

You can delete an existing reallocation scan schedule. However, if you do
this, the job's scan interval reverts to the schedule that was defined for
it when the job was created with the xref:{relative_path}volume-
reallocation-start.html[volume reallocation start] command.

NOTE: This command is not supported for FlexGroup volumes or FlexGroup
constituents.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver.
```

```
``-path <text>`` - Path::
Specifies the path of the reallocation for a LUN, NVMe namespace, file, or
volume.

``[-d, -del <true>]`` - Delete::
Deletes an existing reallocation schedule when set to true. If you use
this command without specifying this parameter, its effective value is
false and the reallocation schedule is not deleted. If you specify this
parameter without a value, it is set to true and the reallocation schedule
is deleted.

``[-s, -cron <text>]`` - Cron Schedule::
Specifies the schedule with the following four fields in sequence. Use a
space between field values. Enclose the values in double quotes.
+
  * minute is a value from 0 to 59.
  * hour is a value from 0 (midnight) to 23 (11:00 p.m.).
  * day of week is a value from 0 (Sunday) to 6 (Saturday).
  * day of month is a value from 1 to 31.
+
NOTE: If you specify 31 as the value for the day of month, reallocation
scans will not run in any months with fewer than 31 days.
+




+
Use an asterisk "*" as a wildcard to indicate every value for that field.
For example, an * in the day of month field means every day of the month.
You cannot use the wildcard in the minute field.
+
You can enter a number, a range, or a comma-separated list of values for a
field.

== Examples
```

cluster1::> volume reallocation schedule -s "0 23 6 *" /vol/db/lun1

```
Schedules a reallocation job to run at 11 pm every Saturday on lun1.

== Related Links

* xref:{relative_path}volume-reallocation-start.html[volume reallocation
start]
```

```
[[ID1e482cc32dc1067aa9f8c1ec2e790d96]]
= volume reallocation show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show reallocate job status

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Displays the status of a reallocation scan, including the state, schedule,
interval, optimization, and log files. If you do not specify the `path`
for a particular reallocation scan, then the command displays all the
reallocation scans.

== 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.

| ``[-v  ]``::
Specify this parameter to display the output in a verbose format.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver::
Specify this parameter to display reallocation scans that match the
Vserver that you specify.

``[-path <text>]`` - Path::
Specify this parameter to display reallocation scans that match the path
that you specify.

``[-t, -threshold <integer>]`` - Threshold::
Specify this parameter to display reallocation scans that match the
```

```
threshold that you specify.

``[-id <integer>]`` - Job ID::
Specify this parameter to display reallocation scans that match the
reallocation job ID that you specify.

``[-description <text>]`` - Job Description::
Specify this parameter to display reallocation scans that match the text
description that you specify.

``[-state
{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Re
schedule|Error|Quit|Dead|Unknown|Restart|Dormant}]`` - Job State::
Specify this parameter to display reallocation jobs that match the state
that you specify.

``[-progress <text>]`` - Execution Progress::
Specify this parameter to list the running reallocation jobs whose
progress indicator matches the text that you provide. For example, if you
specify "Starting ..." as the text string for the progress option, then
the system lists all of the jobs that are starting.

``[-schedule <job_schedule>]`` - Schedule Name::
Specify this parameter to display reallocation scans that match the
schedule name that you specify. If you want a list of all job schedules,
use the xref:{relative_path}job-schedule-show.html[job schedule show]
command.

``[-global-status <text>]`` - Global State of Scans::
Specify this parameter to indicate if reallocation scans are on or off
globally. You must type either of the following text strings:
+
  * "Reallocation scans are on"
  * "Reallocation scans are off"



== Examples
```

cluster1::> volume reallocation show
Vserver Description Schedule State
------- ----------- -------- -----
Reallocation scans are on
vs0 /vol/vol2,space-optimized reallocate_1d Queued

```
Displays the Vserver, description, schedule, and state for the
```

```
reallocation scans on the local node.

== Related Links

* xref:{relative_path}job-schedule-show.html[job schedule show]




[[ID9846e0ca6900b563f54bd96da54d0135]]
= volume reallocation start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start reallocate job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Begins a reallocation scan on a LUN, NVMe namespace, file, or volume when
you specify the path. If a volume has several small files that would
benefit from periodic optimization, specify the /vol/volname.

Before performing a reallocation scan, the reallocation job normally
performs a check of the current layout optimization. If the current layout
optimization is less than the threshold, then the system does not perform
a reallocation on the LUN, NVMe namespace, file, or volume.

You can define the reallocation scan job so that it runs at a specific
interval, or you can use the xref:{relative_path}volume-reallocation-
schedule.html[volume reallocation schedule] command to schedule
reallocation jobs.

NOTE: This command is not supported for FlexGroup volumes or FlexGroup
constituents.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver.
```

```
``-path <text>`` - Path::
Specifies the path of the reallocation for a LUN, NVMe namespace, file, or
volume.

{ ``[-i, -interval <text>]`` - Interval Schedule::
Specifies the reallocation scan interval in
+
  * m for minutes
  * h for hours
  * d for days

+
For example, `30m` is a 30 minute interval. The countdown to the next scan
begins after the first scan is complete.
+
The default interval is 24 hours.

| ``[-o, -once <true>]`` - Once::
Specifies that the job runs once and then is automatically removed from
the system when set to true. If you use this command without specifying
this parameter, its effective value is false and the reallocation scan
runs as scheduled. If you enter this parameter without a value, it is set
to true and a reallocation scan runs once.

| ``[-f, -force <true>]`` - Force }::
Performs a one-time full reallocation on a LUN, file, or volume when set
to true. A forced reallocation rewrites blocks on a LUN, file, or volume
unless the reallocation would result in worse performance. If you use this
command without specifying this parameter, its effective value is false
and a forced reallocation is not performed. If you specify this parameter
without a value, it is set to true, and a forced reallocation is
performed.

{ ``[-p, -space-optimized <true>]`` - Space Optimized::
Specifies that snapshot blocks are not copied to save space when set to
true. If you use this command without specifying this parameter, its
effective value is false and snapshot blocks are copied. However, reads
from snapshots might have a slightly higher latency. If you specify this
parameter without a value, it is set to true and snapshot blocks are not
copied. You cannot use the `space-optimized` option with the `unshare`
option.

| ``[-u, -unshare <true>]`` - Unshare Deduplicated Blocks }::
Specifies that blocks that are shared by deduplication will be unshared.
This option can help remove fragmentation caused on dense volumes. This
```

```
may result in increased disk usage, especially for full reallocation. You
cannot use the `unshare` option with the `space-optimized` option.

{ ``[-t, -threshold <integer>]`` - Threshold::
Specifies the threshold when a LUN, NVMe namespace, file, or volume is
considered unoptimized and a reallocation should be performed. Once the
threshold is reached, the system creates a diagnostic message that
indicates that a reallocation might improve performance.
+
The threshold range is from 3 (the layout is moderately optimized) to 10
(the layout is not optimal). The threshold default is 4.

| ``[-n, -no-check <true>]`` - No Threshold Check }::
Does not check the current layout to determine if a reallocation is needed
when set to true. If you use this command without specifying this
parameter, its effective value is false and the system does check the
current layout to determine if a reallocation is needed. If you specify
this parameter without a value, it is set to true and the system does not
check the current layout to determine if a reallocation is needed.

== Examples
```

cluster1::> volume reallocation start -path /vol/vol2 -interval 30m
[Job 165] Job is queued: Reallocate Job.

```
Starts a reallocation job on volume vol2 every 30 minutes.

== Related Links

* xref:{relative_path}volume-reallocation-schedule.html[volume
reallocation schedule]




[[ID48e4952c9f247efab52476975907308e]]
= volume reallocation stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop reallocate job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Stops and deletes any reallocation scans on a LUN, NVMe namespace, file,
or volume. This command stops and deletes in-progress, scheduled, and
quiesced scans.

NOTE: This command is not supported for FlexGroup volumes or FlexGroup
constituents.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver.

``-path <text>`` - Path::
Specifies the path of the reallocation for a LUN, NVMe namespace, file, or
volume.

== Examples
```

cluster1::> volume reallocation stop /vol/vol2
1 entry was deleted.

```
Stops and deletes one reallocation scan on volume vol2.




:leveloffset: -1


= volume rebalance commands

:leveloffset: +1


[[ID3643438653f85694f34f416818af9912]]
= volume rebalance modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the configuration for volume capacity rebalancing operations.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume rebalance modify` command modifies the configuration for
volume capacity rebalancing operations on a FlexGroup volume in the
specified Vserver. When a new capacity rebalancing operation is started on
a FlexGroup volume, it will use the current configuration values. Once the
operation is started, any changes to the configuration will not affect the
currently running capacity rebalancing operation. Only future capacity
rebalancing operations will use the new configuration values.

The `volume rebalance modify` command is only supported on FlexGroup
volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Identifies the Vserver that contains the FlexGroup volume whose capacity
```

```
rebalancing configuration will be modified.

``-volume <volume name>`` - Volume Name::
Identifies the FlexGroup volume whose capacity rebalancing configuration
will be modified.

``[-max-runtime <time_interval>]`` - Maximum Runtime::
Specifies the maximum time a capacity rebalancing operation will run for.

``[-max-threshold <percent>]`` - Maximum Imbalance Threshold per
Constituent::
Specifies the maximum imbalance percentage for constituents. When a
constituent's imbalance percentage is larger than this value, files will
be moved from the constituent. The default value is 20%.

``[-min-threshold <percent>]`` - Minimum Imbalance Threshold per
Constituent::
Specifies the minimum imbalance percentage for constituents. When a
constituent's imbalance percentage is smaller than this value, files will
not be moved from the constituent. The default value is 5%.

``[-max-file-moves <integer>]`` - Maximum Concurrent File Moves per
Constituent::
Specifies the maximum number of concurrent file moves in a volume capacity
rebalancing operation on a constituent of the FlexGroup volume. The
default value is 25.

``[-min-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum File Size::
Specifies the minimum file size to consider for a volume capacity
rebalancing operation. The value must be a multiple of 4KB. The default
value is 100MB, and the minimum configurable value is 20MB. Setting the
minimum file size to less than the default value leads to more files being
moved. Moved files use granular data, which might impact read/write I/O
performance.

``[-exclude-snapshots {true|false}]`` - Exclude Files Stuck in Snapshot
Copies::
Specifies whether files stuck in snapshots should be excluded in a volume
capacity rebalancing operation. The default value is true.

``[-start-time <Date>]`` - Scheduled start time::
Specifies the scheduled start time of the volume capacity rebalancing
operation. The value must be a time later than 'now'.

== Examples

In this example, the volume capacity rebalancing configuration is modified
```

```
for the FlexGroup volume "fg1" in vserver "vs0" with the maximum threshold
being 30%, the minimum threshold being 10%, the maximum file moves being
35, and the minimum file size being 200MB.
```

cluster::volume rebalance*> modify -vserver vs0 -volume fg1 -max-threshold 30 -min-threshold 10 -max-file -moves 35 -min-file-size 200MB

```
In this example, a previously scheduled volume capacity rebalancing is
rescheduled for a different time.
```

cluster::volume rebalance*> modify -vserver vs0 -volume fg1 -start-time 12/30/2022 22:59:59

```
[[ID2dbbd730266df72edb0d3dc2166675f0]]
= volume rebalance show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of volume capacity rebalancing operations on FlexGroup
volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume rebalance show` command displays information about volume
capacity rebalancing operations on FlexGroup volumes.

The `volume rebalance show` command is only supported on FlexGroup volumes
and constituents. The command displays rebalancing information for
volumes. The command output depends on the parameter or parameters
specified with the command. If no parameters are specified, the command
displays the following information about all volumes:

* Vserver name
* Volume name
* Rebalancing state
* Total size
* Used size for calculating imbalance
* Size used by snapshots
* Target used size
* Imbalance size
```

```
* Imbalance percent



To display detailed rebalancing information about each FlexGroup volume or
constituent, run the command with the `-instance` parameter. The detailed
view provides the following information for FlexGroup volumes:

* Vserver name
* Volume name
* Is constituent
* Rebalance state
* Total size
* Used size for calculating imbalance
* Size used by snapshots
* Imbalance size
* Imbalance percentage
* Maximum constituent imbalance percentage
* Rebalance start time
* Rebalance stop time
* Rebalance runtime
* Rebalance maximum runtime
* Maximum imbalance threshold per constituent
* Minimum imbalance threshold per constituent
* Maximum files moved per constituent
* Minimum file size
* Exclude files stuck in snapshots
* Rebalance notice messages



The detailed view provides the following information for FlexGroup
constituents:

* Vserver name
* Volume name
* Is constituent
* Rebalance state
* Total size
* Used size for calculating imbalance
* Size used by snapshots
* Constituent target used size
* Imbalance size
* Imbalance percentage
* Moved data size
* Rebalance start time
```

```
* Rebalance stop time
* Rebalance runtime
* Rebalance maximum runtime
* Rebalance maximum runtime
* Maximum imbalance threshold per constituent
* Minimum imbalance threshold per constituent
* Maximum files moved per constituent
* Minimum file size
* Exclude files stuck in snapshots



You can specify parameters from the above list to display information. For
example, to display information only about currently running rebalancing
operations, run the command with the `-state rebalancing` parameter

== 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.

| ``[-config  ]``::
If this parameter is specified, the command displays the following
information:
+
  * max-runtime
  * max-threshold
  * min-threshold
  * max-file-moves
  * min-file-size
  * exclude-snapshots



| ``[-time  ]``::
If this parameter is specified, the command displays the following
information:
+
  * start-time
  * runtime
  * max-runtime



| ``[-instance  ]`` }::
```

```
If this parameter is specified, the command displays the additional
detailed information described above.

``[-vserver <vserver name>]`` - Vserver Name::
Selects information about volume capacity rebalancing operations on
FlexGroup volumes in the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
Selects information about volume capacity rebalancing operations on the
specified FlexGroup volume.

``[-is-constituent {true|false}]`` - Is Constituent::
If specified, the command displays information about volume capacity
rebalancing operations on constituents.

``[-state {idle|starting|rebalancing|rebalancing-source|rebalancing-
dest|scanning|stopping|paused|not-running|unknown|scheduled}]`` -
Rebalance State::
Selects information about volume capacity rebalancing operations with the
specified state. While a FlexGroup volume is rebalancing, every
constituent will have a rebalancing engine that can either be scanning the
filesystem for space usage and files to move, actively moving files or
temporarily doing neither. If one or more constituents has a state of
``_rebalancing-source_`` or ``_rebalancing-dest_`` , then files are being
moved to rebalance the FlexGroup. If no files are being moved, more
information about what the rebalancing engine is doing for each
constituent is available using the (privilege: advanced) `volume rebalance
show -instance` command. The following values apply to FlexGroup volumes.
+
  * ``_not-running_`` - Capacity rebalancing is not running on the volume.
This will only change when the user explicitly starts rebalancing.
  * ``_idle_`` - Capacity rebalancing is running, however, no active
scanning or file movement is currently happening.
  * ``_starting_`` - The capacity rebalancing operation is starting on the
volume.
  * ``_rebalancing_`` - Volume capacity rebalancing is running on the
volume.
  * ``_paused_`` - The capacity rebalancing operation is paused on the
volume.
  * ``_stopping_`` - The capacity rebalancing operation is stopping on the
volume.
  * ``_unknown_`` - The system was unable to determine the rebalancing
state for the volume.
```

```
``[-notices <text>,...]`` - Rebalance Notice Messages::
Selects information about volume capacity rebalancing operations with the
specified notice messages.

``[-total {<integer>[KB|MB|GB|TB|PB]}]`` - Total Size::
Selects information about volume capacity rebalancing operations with the
specified total size.

``[-used-for-imbalance {<integer>[KB|MB|GB|TB|PB]}]`` - Used Size for
Calculating Imbalance::
Selects information about volume capacity rebalancing operations with the
specified used size for calculating imbalance. This is calculated by
subtracting the size used by snapshots, the size of files pending deletion
and the size used by filesystem metadata from the volume's used size.

``[-size-used-by-snapshots {<integer>[KB|MB|GB|TB|PB]}]`` - Volume Size
Used by Snapshot Copies::
Selects information about volume capacity rebalancing operations with the
specified volume size used by snapshots.

``[-target-used {<integer>[KB|MB|GB|TB|PB]}]`` - Constituent Target Used
Size::
Selects information about volume capacity rebalancing operations with the
specified target used size on a constituent. Represents the average used
size among all constituents.

``[-imbalance-size {<integer>[KB|MB|GB|TB|PB]}]`` - Imbalance Size::
Selects information about volume capacity rebalancing operations with the
specified imbalance size.

``[-imbalance-percent <percent_no_limit_signed>]`` - Imbalance
Percentage::
Selects information about volume capacity rebalancing operations with the
specified imbalance percentage.

``[-data-moved {<integer>[KB|MB|GB|TB|PB]}]`` - Moved Data Size::
Selects information about volume capacity rebalancing operations with the
specified moved data size.

``[-max-constituent-imbalance-percent <percent_no_limit>]`` - Maximum
Constituent Imbalance Percentage::
Selects information about volume capacity rebalancing operations with the
specified maximum constituent imbalance percentage. Represents the
imbalance percentage of the constituent most out of balance.

``[-start-time <Date>]`` - Rebalance Start Time::
```

```
Selects information about volume capacity rebalancing operations with the
specified start time.

``[-stop-time <Date>]`` - Rebalance Stop Time::
Selects information about volume capacity rebalancing operations with the
specified stop time.

``[-runtime <time_interval>]`` - Rebalance Runtime::
Selects information about volume capacity rebalancing operations with the
specified run time.

``[-max-runtime <time_interval>]`` - Rebalance Maximum Runtime::
Selects information about volume capacity rebalancing operations with the
specified maximum run time. Represents the maximum amount of time a
rebalancing operation can run for before stopping.

``[-max-threshold <percent>]`` - Maximum Imbalance Threshold per
Constituent::
Selects information about volume capacity rebalancing operations with the
specified maximum threshold. Represents the acceptable imbalance threshold
beyond which rebalancing should be started on that constituent in the
FlexGroup volume.

``[-min-threshold <percent>]`` - Minimum Imbalance Threshold per
Constituent::
Selects information about volume capacity rebalancing operations with the
specified minimum threshold. Represents the minimum imbalance threshold at
which rebalancing is stopped for the constituent in the FlexGroup volume.

``[-max-file-moves <integer>]`` - Maximum Concurrent File Moves per
Constituent::
Selects information about volume capacity rebalancing operations with the
specified maximum concurrent file moves. Represents the maximum number of
concurrent file moves allowed per constituent.

``[-min-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum File Size::
Selects information about volume capacity rebalancing operations with the
specified minimum file size. Represents the minimum size of file to be
considered for rebalancing a constituent in a FlexGroup volume.

``[-exclude-snapshots {true|false}]`` - Exclude Files Stuck in Snapshot
Copies::
Selects information about volume capacity rebalancing operations with the
specified excluding snapshot flag. Represents whether or not files that
are in snapshots as candidates for rebalancing are ignored.

== Examples
```

```
The following example displays rebalancing information about all FlexGroup
volumes on Vserver "vs0":
```

cluster::> volume rebalance show -vserver vs0
Vserver: vs0
Projected
AFS Target Imbalance
Volume State Total Used Used Size %
---------- ---------- --------- --------- --------- --------- -----
fg1 rebalancing 400MB 117.7MB - 176KB 0%
fg2 not-running 800MB 115.0MB - 8KB 0%
2 entries were displayed.

```
The following example displays rebalancing information about all FlexGroup
constituents on Vserver "vs0":
```

cluster::> volume rebalance show -vserver vs0 -is-constituent True
Vserver: vs0
Projected
AFS Target Imbalance
Volume State Total Used Used Size %
---------- ---------- --------- --------- --------- --------- -----
fg1*0001 idle 200MB 59.03MB 58.86MB 176KB 0%*
*fg1*0002 idle 200MB 58.69MB 58.86MB -172KB 0%
fg2*0001 not-running 200MB 28.77MB 28.76MB 4KB 0%*
*fg2*0002 not-running 200MB 28.76MB 28.76MB -4KB 0%
fg2*0003 not-running 200MB 28.77MB 28.76MB 4KB 0%*
*fg2*0004 not-running 200MB 28.76MB 28.76MB -4KB 0%
6 entries were displayed.

```
The following example displays detailed rebalancing information about a
specific FlexGroup volume "fg1" on Vserver "vs0":
```

cluster::> volume rebalance show -vserver vs0 -volume fg1 -instance
Vserver Name: vs0
Volume Name: fg1
Is Constituent: false
Rebalance State: not-running
Rebalance Notice Messages: -
Total Size: 400MB
Projected AFS Used Size: 117.7MB
Volume Size Used by Snapshot Copies: 10.57MB
Constituent Target Used Size: -
Imbalance Size: 176KB
Imbalance Percentage: 0%
Moved Data Size: -
Maximum Constituent Imbalance Percentage: 0%
Rebalance Start Time: Fri Mar 25 17:26:09 2022

Rebalance Stop Time: Fri Mar 25 17:29:29 2022
Rebalance Runtime: 0h3m20s
Rebalance Maximum Runtime: 0h3m20s
Maximum Imbalance Threshold per Constituent: 20%
Minimum Imbalance Threshold per Constituent: 5%
Maximum Concurrent File Moves per Constituent: 25
Minimum File Size: 4KB
Exclude Files Stuck in Snapshot Copies: true

```
The following example displays detailed rebalancing information about a
specific FlexGroup constituent "fg1__0001" on Vserver "vs0":
```

cluster::> volume rebalance show -vserver vs0 -volume fg1*0001 -instance -is-constituent True*
Vserver Name: vs0
*Volume Name: fg1*0001
Is Constituent: true
Rebalance State: not-running
Rebalance Notice Messages: -
Total Size: 200MB
Projected AFS Used Size: 59.04MB
Volume Size Used by Snapshot Copies: 5.46MB
Constituent Target Used Size: 58.86MB
Imbalance Size: 176KB
Imbalance Percentage: 0%
Moved Data Size: 0B
Maximum Constituent Imbalance Percentage: -
Rebalance Start Time: Fri Mar 25 17:26:09 2022
Rebalance Stop Time: Fri Mar 25 17:29:29 2022
Rebalance Runtime: 0h3m20s
Rebalance Maximum Runtime: 0h3m20s
Maximum Imbalance Threshold per Constituent: 20%
Minimum Imbalance Threshold per Constituent: 5%
Maximum Concurrent File Moves per Constituent: 25
Minimum File Size: 4KB
Exclude Files Stuck in Snapshot Copies: true

```
[[ID52c50fb42c3c73adebc7f8da53c4e4e4]]
= volume rebalance start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a volume capacity rebalancing operation on a FlexGroup volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `volume rebalance start` command starts a volume capacity rebalancing
operation on a FlexGroup volume in the specified Vserver. Only one volume
capacity rebalancing operation is supported per FlexGroup volume at a
time. Running capacity rebalancing operations on multiple FlexGroup
volumes at the same time is supported.

The `volume rebalance start` command is only supported on FlexGroup
volumes. If granular data is not already enabled on the volume, it will
automatically be set to ``_basic_`` when this command is run.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Identifies the Vserver that contains the FlexGroup volume whose capacity
will be rebalanced.

``-volume <volume name>`` - Volume Name::
Identifies the FlexGroup volume whose capacity will be rebalanced.

``[-max-runtime <time_interval>]`` - Maximum Runtime::
Specifies the maximum time a capacity rebalancing operation will run for.
The default value is 6 hours.

``[-max-threshold <percent>]`` - Maximum Imbalance Threshold per
Constituent::
Specifies the maximum imbalance percentage for constituents. When a
constituent's imbalance percentage is larger than this value, files will
be moved from the constituent. The default value is 20%.

``[-min-threshold <percent>]`` - Minimum Imbalance Threshold per
Constituent::
Specifies the minimum imbalance percentage for constituents. When a
constituent's imbalance percentage is smaller than this value, files will
not be moved from the constituent. The default value is 5%.

``[-max-file-moves <integer>]`` - Maximum Concurrent File Moves per
Constituent::
Specifies the maximum number of concurrent file moves in a volume capacity
rebalancing operation on a constituent of the FlexGroup volume. The
default value is 25.

``[-min-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Minimum File Size::
Specifies the minimum file size to consider for a volume capacity
rebalancing operation. The value must be a multiple of 4KB. The default
value is 100MB, and the minimum configurable value is 20MB. Setting the
```

```
minimum file size to less than the default value leads to more files being
moved. Moved files use granular data, which might impact read/write I/O
performance.

``[-exclude-snapshots {true|false}]`` - Exclude Files Stuck in Snapshot
Copies::
Specifies whether files stuck in snapshots should be excluded in a volume
capacity rebalancing operation. The default value is true.

``[-start-time <Date>]`` - Scheduled start time, e.g. 12/31/2022
23:59:59::
Specifies the scheduled start time of the volume capacity rebalancing
operation. The value must be a time later than 'now'.

== Examples

In this example, a volume capacity rebalancing operation is started for
the FlexGroup volume "fg1" in vserver "vs0" with the maximum run time
being 72000 seconds or 20 hours starting at 23:59:59 on 12/31/2022.
```

cluster::volume rebalance*> start -vserver vs0 -volume fg1 -max-runtime 72000 -start-time 12/31/2022 23:59:59

```
[[ID3448031a31de94c7a24a97b45bbb5111]]
= volume rebalance stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop a volume capacity rebalancing operation on a FlexGroup volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume rebalance stop` command stops a volume capacity rebalancing
operation on a FlexGroup volume in the specified Vserver.

The `volume rebalance stop` command is only supported on FlexGroup
volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Identifies the Vserver that contains the FlexGroup volume whose capacity
rebalancing operation will be stopped.

``-volume <volume name>`` - Volume Name::
Identifies the FlexGroup volume whose capacity rebalancing operation will
be stopped.

== Examples

In this example, a volume capacity rebalancing operation is stopped for
the FlexGroup volume "fg1" in vserver "vs0".
```

cluster::volume rebalance rebalance*> stop -vserver vs0 -volume fg1

```
[[ID986a97f93e0aa97c747de84dc1fd89a4]]
= volume rebalance file-move abort
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Abort a file-move operation that is in progress

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume rebalance file-move abort` command aborts a file-move
operation that has been started and is in progress.

A file-move operation goes through different stages as part of the move
operation. During the cutover stage of the file-move operation, the actual
file is created at the destination FlexGroup constituent, after which the
data blocks are transferred. Once the file is created at the destination
FlexGroup constituent, the file is visible in the namespace of the
FlexGroup volume. If the file-move operation has entered the cutover
stage, it cannot be aborted.

The `volume rebalance file-move abort` command is only supported on
FlexGroup volumes. For a given FlexGroup volume, there can be many file-
move operations that are currently in progress for different sets of
constituents. The parameters for this command provide a way to identify a
particular file-move operation within a FlexGroup volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Identifies the Vserver that contains the FlexGroup volume involved in the
file-move operation.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
Identifies the FlexGroup volume that is hosting the file being moved in
the file-move operation.

``-path <text>`` - Path [.advaced-priv]#(privilege: advanced)#::
Identifies the path to the file being moved within the FlexGroup volume.

``-destination-constituent <volume name>`` - Destination Constituent
Volume [.advaced-priv]#(privilege: advanced)#::
Identifies the destination constituent name of the FlexGroup volume where
the file is being moved.

``[-operation-uuid <UUID>]`` - Operation UUID [.advaced-priv]#(privilege:
advanced)#::
```

```
If specified, this parameter identifies the UUID of the file-move
operation. Use the xref:{relative_path}volume-rebalance-file-move-
show.html[volume rebalance file-move show] command to obtain the UUID for
a file-move operation.

== Examples

The following example aborts a file-move operation for the file
"/system/build/image.tgz" within the FlexGroup volume "bld_volume":
```

cluster1::*> volume rebalance file-move abort -vserver vs0 -volume bld_volume -path /system/build/image.tgz -destination-constituent bld_volume__0005

File-move operation cancelled.

```
The following example aborts the file-move operation in progress for the
file "/system/build/image.tgz" within the FlexGroup "bld_volume". The UUID
of the file-move operation is also provided:
```

cluster1::*> volume rebalance file-move abort -vserver vs0 -volume bld_volume -path /system/build/image.tgz -destination-constituent bld_volume__0005 -operation-uuid f8d780b4-32fc-4053-be5e-2f6edc0a652b

File-move operation cancelled.

```
The following example attempts to abort the file-move operation for the
file "/system/build/image.tgz" within the FlexGroup volume "bld_volume".
However, the file-move operation could not be aborted as it is past the
cutover stage of the operation. A file-move operation can only be aborted
if it has not entered the cutover stage:
```

cluster1::*> volume rebalance file-move abort -vserver vs0 -volume bld_volume -path /system/build/image.tgz -destination-constituent bld_volume__0005

Error: command failed: The file move abort operation failed. Reason: Cannot destroy file operation. File move operation has completed the cutover phase.

```
== Related Links

* xref:{relative_path}volume-rebalance-file-move-show.html[volume
rebalance file-move show]




[[ID4243765dcc2c6e836422ff889f89f415]]
```

```
= volume rebalance file-move modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a file-move operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume rebalance file-move modify` command modifies certain
attributes of a file-move operation that is currently in progress.
Specifically, the maximum throughput allowed on a particular file-move
operation can be modified, and the data transfer of a file-move operation
can be paused or resumed.

The `volume rebalance file-move modify` command is only supported on
FlexGroup volumes. For a given FlexGroup volume, there can be many file-
move operations that are currently in progress for different sets of
constituents. The parameters for this command provide a way to identify a
particular file-move operation within a FlexGroup volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Identifies the Vserver that contains the FlexGroup volume involved in the
file-move operation.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
Identifies the FlexGroup volume that is hosting the file being moved in
the file-move operation.

``-path <text>`` - Path [.advaced-priv]#(privilege: advanced)#::
Identifies the path to the file being moved within the FlexGroup volume.
This parameter can be specified with the `-destination-constituent`
parameter to identify a particular file-move operation within a FlexGroup
volume.

``-destination-constituent <volume name>`` - Destination Constituent
Volume [.advaced-priv]#(privilege: advanced)#::
Identifies the destination constituent name of the FlexGroup volume where
```

```
the file is being moved. This parameter can be specified with the `-path`
parameter to identify a particular file-move operation within a FlexGroup
volume.

``[-operation-uuid <UUID>]`` - Operation UUID [.advaced-priv]#(privilege:
advanced)#::
Identifies the UUID of the file-move operation. Use the
xref:{relative_path}volume-rebalance-file-move-show.html[volume rebalance
file-move show] command to obtain the UUID for a file-move operation. This
parameter can be specified in addition to the `-path` and `-destination-
constituent` parameters to identify a particular file-move operation if
there are multiple ongoing operations involving the same path and
destination constituent.

``[-max-throughput {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Scanner Speed
[.advaced-priv]#(privilege: advanced)#::
If specified, this parameter will modify the rate limit for the data
transfer in bytes per second. A non-zero value less than 1 MB/s will be
set to 1 MB/s. A non-zero value greater than 1 MB/s will be truncated to
the nearest integral megabyte value. A value of "0" specifies that no
range will be set for the data transfer. This value does not represent a
throughput the system will guarantee, but a limit that cannot be exceeded.

``[-scanner-paused {true|false}]`` - Scanner Paused [.advaced-
priv]#(privilege: advanced)#::
If specified, the data transfer of the file-move operation will be paused
or resumed. Using a value of ``_false_`` will pause the data transfer
while using a value of ``_true_`` will resume the data transfer.

== Examples

The following example modifies the maximum throughput of the ongoing file-
move operation for the file "/system/build/image.tgz" within the FlexGroup
volume "bld_volume". `-path` and `-destination-constituent` are used to
identify the particular file-move operation:
```

cluster1::*> volume rebalance file-move modify -vserver vs0 -volume bld_volume -path /system/build/image.tgz -destination-constituent bld_volume__0005 -max-throughput 2MB

```
The following example pauses the data transfer of the ongoing file-move
operation for the file "/system/build/image.tgz" within the FlexGroup
volume "bld_volume". This stops any further data from being transferred
from the source volume to the destination volume:
```

cluster1::*> volume rebalance file-move modify -vserver vs0 -volume bld_volume -path /system/build/image.tgz

-destination-constituent bld_volume__0005 -scanner-paused true

```
The following example resumes the data transfer of a paused file-move
operation for the file "/system/build/image.tgz" within the FlexGroup
volume "bld_volume". `-path` and `-destination-constituent` are used to
identify the particular file-move operation:
```

cluster1::*> volume rebalance file-move modify -vserver vs0 -volume bld_volume -path /system/build/image.tgz -destination-constituent bld_volume__0005 -scanner-paused false

```
== Related Links

* xref:{relative_path}volume-rebalance-file-move-show.html[volume
rebalance file-move show]




[[IDf316e46ad53591a6a5667f1904e854c9]]
= volume rebalance file-move show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of files being moved

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume rebalance file-move show` command displays a list of files
being moved between the FlexGroup Constituents.

The `volume rebalance file-move show` command is only supported on
FlexGroup volumes. The command displays information about each ongoing
file-move operation, as well as any file-move operations that have
completed, failed or have been cancelled within the last two minutes. For
all file-move operations matching the specified parameters, the command
displays the following information:

* Vserver name
* Volume name
* File Path
```

```
* Source Constituent Volume
* Destination Constituent Volume
* Current status of the file-move operation ("Preparing", "Data-Transfer",
"Paused", "Complete", "Failed" or "Terminated-Manual")
* Percentage of total data bytes transferred to the destination



To display detailed information about each file-move operation, run the
command with the `-instance` parameter. The detailed view provides the
following information:

* Vserver Name
* Volume Name
* Destination Constituent Volume
* Source Constituent Volume
* Destination Constituent Node
* Source Constituent Node
* Operation Status: Current status of the file-move operation
("Preparing", "Data-Transfer", "Paused", "Complete", "Failed" or
"Terminated-Manual")
* Scanner Progress Percentage
* Operation UUID: UUID of the overall file-move operation
* Source Operation UUID: UUID used to identify the file-move operation on
the source constituent
* File Path: Path of the file being moved
* Source Constituent Master Data Set ID
* Source Constituent Data Set ID
* Source File ID: ID of the inode on the source constituent
* Source Generation: Generation number of the inode on the source
constituent
* Destination File ID: ID of the inode on the destination constituent
* Destination Data Set ID
* Destination Master Data Set ID
* Destination Generation: Generation number of the inode on the
destination constituent
* Maximum Scanner Speed (per sec): Maximum throughput of the data scanner
* Scanner Paused: Displays information if the underlying data transfer
scanner is paused
* Scanner Status: The current stage of the underlying data transfer
scanner ("Preparing", "Allocation-Map", "Data", "Destroying", "Paused-
Manual", "Paused-Error", "Complete", "Terminated-Manual" or "Destroyed")
* Scanner Progress: Number of bytes transferred
* Scanner Total: Total number of bytes to be transferred
* Data Scanner Priority: Priority of the data scanner assigned to the
file-move operation (high or low)
```

```
* Elapsed Time
* Cutover Time: Time elapsed during the cutover stage of the file-move
operation, in seconds
* Is Snapshot Fenced: Whether snapshots are fenced for the file-move
operation
* Is Destination Ready: Whether the destination volume is ready for the
file-move operation
* Last Failure Time
* Last Failure Reason



You can specify parameters from the above list to display information. For
example, to display information only about paused operations, run the
command with the `-status Paused` parameter.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is specified, the command displays the additional
detailed information described above.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-volume` parameter are specified, the command
displays the file-move operations on the specified volume. If this
parameter is specified by itself, the command displays the file-move
operations on all volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-vserver` parameter are specified, the command
displays the file-move operations on the specified volume. If this
parameter is specified by itself, the command displays the file-move
operations on all volumes matching the specified name.

``[-destination-constituent <volume name>]`` - Destination Constituent
Volume [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only the file-move
operations with the specified destination FlexGroup constituent name.

``[-operation-uuid <UUID>]`` - Operation UUID [.advaced-priv]#(privilege:
```

```
advanced)#::
If this parameter is specified, the command displays only the file-move
operation with the specified operation UUID.

``[-destination-node <nodename>]`` - Destination Constituent Node
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations for files being moved to the specified destination node.

``[-destination-dsid <integer>]`` - Destination Constituent DSID
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations with the specified destination volume data set ID.

``[-source-constituent <volume name>]`` - Source Constituent Volume
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations with the specified source FlexGroup constituent name.

``[-source-operation-uuid <UUID>]`` - Source Operation UUID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only the file-move
operation with the specified internal source job operation UUID.

``[-source-node <nodename>]`` - Source Constituent Node [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations for files being moved from the specified source node.

``[-source-dsid <integer>]`` - Source Constituent DSID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations with the specified source volume data set ID.

``[-status {Preparing|Data-Transfer|Paused|Complete|Failed|Manually-
Terminated}]`` - Operation Status [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations with the specified status ("Preparing", "Data-Transfer",
"Paused", "Complete", "Failed" or "Terminated-Manual").

``[-scanner-percent <percent>]`` - Scanner Progress Percentage [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations that have transferred the specified percentage of total bytes.

``[-last-failure-reason <text>]`` - Last Failure Reason [.advaced-
```

```
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations whose last failure reason matches the specified string.

``[-path <text>]`` - File Path [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations for files that match the specified path.

``[-destination-msid <integer>]`` - Destination Constituent MSID
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations with the specified destination volume master data set ID.

``[-destination-fileid <integer>]`` - Destination Fileid [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations pertaining to files with the specified file ID on the
destination constituent.

``[-destination-generation <integer>]`` - Destination Generation
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations pertaining to files with the specified generation number on the
destination constituent.

``[-source-msid <integer>]`` - Source Constituent MSID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations with the specified source volume master data set ID.

``[-source-fileid <integer>]`` - Source Fileid [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays only file-move
operations pertaining to files with the specified file ID on the source
constituent.

``[-source-generation <integer>]`` - Source Generation [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations pertaining to files with the specified generation number on the
source constituent.

``[-max-throughput {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Scanner Speed
(per sec) [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations that use the specified value as the maximum scanner speed (in
```

```
MB/s). This value does not represent a throughput the system will
guarantee, but a limit that cannot be exceeded.

``[-scanner-paused {true|false}]`` - Scanner Paused [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations that are paused (true) or unpaused (false).

``[-scanner-status {Preparing|Allocation-Map|Data|Destroying|Paused-
Manual|Paused-Error|Complete|Destroyed|Terminated-Manual}]`` - Scanner
Status [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations in the specified scanner stage ("Preparing", "Allocation-Map",
"Data", "Destroying", "Paused-Manual", "Paused-Error", "Complete",
"Terminated-Manual" or "Destroyed")

``[-scanner-progress {<integer>[KB|MB|GB|TB|PB]}]`` - Scanner Progress
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations that have transferred the specified number of bytes.

``[-scanner-total {<integer>[KB|MB|GB|TB|PB]}]`` - Scanner Total
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations with the specified total transfer size in bytes.

``[-data-scanner-priority {high|low}]`` - Data Scanner Priority [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only the file-move
operations that give the data scanner high priority (high) or low priority
(low).

``[-elapsed-time <time_interval>]`` - Elapsed Time [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only file-move
operations that have a matching `elapsed-time` interval. Specify an
`elapsed-time` range by using the ".." operator between two values. For
example, the following command displays file-move operations with
`elapsed-time` between 2 hours and 3 hours 30 minutes:
```

volume rebalance file-move show -elapsed-time 2h..3h30m

```
+
Specify a comparative `elapsed-time` value using the ">" and "<"
operators. For example, the following command displays the file-move
operations with an `elapsed-time` greater than 5 hours and 30 minutes:
```

volume rebalance file-move show -elapsed-time >5h30m

```
``[-cutover-time <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Cutover Time [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only the file-move
operations with the specified cutover time.

``[-is-snapshot-fenced {true|false}]`` - Is Snapshot Fenced [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only the file-move
operations with snapshots fenced (true) or unfenced (false).

``[-is-destination-ready {true|false}]`` - Is Destination Ready [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays only the file-move
operations for which the destination volume is ready (true) or not ready
(false).

``[-last-failure-time <MM/DD/YYYY HH:MM:SS>]`` - Last Failure Time
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays file-move operations
that have a matching `last-failure-time` value. Use the format:
"MM/DD/YYYY HH:MM:SS". Specify a time range by using the ".." operator
between two time statements. For example, the following command displays
file-move operations with a `last-failure-time` value between "08/13/2014
05:55:00" and "08/13/2014 06:10:00":
```

volume rebalance file-move show -last-failure-time "08/13/2021 05:55:00".."08/13/2021 06:10:00"

```
+
Specify a comparative `last-failure-time` value using the ">" and "<"
operators. For example, the following command displays file-move
operations with a `last-failure-time` value after "8/27/2021 14:58:00":
```

volume rebalance file-move show -last-failure-time >"8/27/2021 14:58:00"

```
== Examples

The following example displays information about all file-move operations
on the Vserver named "vs1":
```

cluster1::volume rebalance file-move*> show -vserver vs1
Vserver: vs1
Volume: fg1
Path Source Destination Status Progress
---------- ---------- ---------- ---------- ---------
/d1/d2/foo fg1*0002 fg1*0003 Paused 0%
/d1/d2/bar fg1*0002 fg1*0004 Completed 100%

Vserver: vs1
Volume: bld_volume
Path Source Destination Status Progress
---------- ---------- ---------- ---------- ---------
/system/build/image.tgz bld_volume*0001*
*bld_volume*0005
Data-Transfer 50%
3 entries were displayed.

```
The following example displays detailed information about a specific file-
move operation on Vserver vs0:
```

cluster1::volume rebalance file-move*> show -instance -vserver vs0 -volume fg -path /test
(volume rebalance file-move show)
Vserver Name: vs0
Volume Name: fg
Destination Constituent Volume: fg*0001*
Operation UUID: 283a983e-c06a-42bc-8bd7-8830211cb2a9
Destination Constituent Node: sti96-vsim-ucs539i
Destination Constituent DSID: 1635
*Source Constituent Volume: fg*0004
Source Operation UUID: 283a983e-c06a-42bc-8bd7-8830211cb2a9
Source Constituent Node: sti96-vsim-ucs539j
Source Constituent DSID: 1638
Operation Status: Paused
Scanner Progress Percentage: 0%
Last Failure Reason: -
File Path: /test
Destination Constituent MSID: 2154447497
Destination Fileid: 97
Destination Generation: 247243720
Source Constituent MSID: 2154447500
Source Fileid: 504
Source Generation: 214935427
Maximum Scanner Speed (per sec): 1MB
Scanner Paused: true
Scanner Status: Paused-Manual

Scanner Progress: 0B
Scanner Total: 10MB
Data Scanner Priority: low
Elapsed Time: 0h2m27s
Cutover Time: 0s
Is Snapshot Fenced: false
Is Destination Ready: true
Last Failure Time: -

```
[[ID35b469b368e05d70b5ebac2380c9ecf1]]
= volume rebalance file-move start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Moves a file from one FlexGroup constituent to another

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume rebalance file-move start` command initiates a retroactive
file movement for a file residing in a FlexGroup volume in the given
Vserver. You must provide the path to the file to be moved. Additional
parameters may be used to control certain properties of the move
operation. The parent directory's mtime may be modified by this command
which may impact analytics mtime histogram. This command only takes NFS
filenames and only UTF-8 encoding is supported.

The `volume rebalance file-move start` command is only supported on
FlexGroup volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Identifies the Vserver that contains the FlexGroup volume involved in this
operation.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
Identifies the FlexGroup volume that is hosting the file being moved in
this operation.
```

```
``-path <text>`` - Path [.advaced-priv]#(privilege: advanced)#::
Path to the file being moved within this FlexGroup volume.

``[-destination-constituent <volume name>]`` - Destination Constituent
Volume [.advaced-priv]#(privilege: advanced)#::
Identifies the destination constituent name of the FlexGroup volume where
the file is being moved to. The constituent must be a member of the same
FlexGroup volume as the source constituent.

``[-max-throughput {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Scanner Speed
[.advaced-priv]#(privilege: advanced)#::
Specifies the range of the data transfer in bytes per second. A non-zero
value less than 1 MB/s will be set to 1 MB/s. A non-zero value greater
than 1 MB/s will be truncated to the nearest integral megabyte value. If
unspecified, the default value is "0" which means no range will be set for
the data transfer. This value does not represent a throughput the system
will guarantee, but a limit that cannot be exceeded.

``[-scanner-paused {true|false}]`` - Scanner Paused [.advaced-
priv]#(privilege: advanced)#::
Specifies that the data scanner should not be started immediately when the
move operations starts. If unspecified, the default value is "false". Use
the xref:{relative_path}volume-rebalance-file-move-modify.html[volume
rebalance file-move modify] command to start the data scanner.

``[-force <true>]`` - Force [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the file move operation will break the
existing lock state on the file and proceed. This may cause a disruption
for some client applications.

``[-is-disruptive {true|false}]`` - Is Disruptive [.advaced-
priv]#(privilege: advanced)#::
Specifies whether the file move operation should be disruptive. The
default value is "false".

== Examples

In this example, a file-move operation is started for the file
"/d1/d2/foo" within the FlexGroup volume "fg1" in vserver "vs0". The path
is relative to the root directory of this FlexGroup volume, not the global
namespace path of the file. In this example, the file is being moved to
the destination constituent identified by the name "fg1__0003".
```

cluster::volume rebalance file-move*> start -vserver vs0 -volume fg1 -path /d1/d2/foo -destination-constituent fg1__0003
File move started with operation-uuid f8d780b4-32fc-4053-be5e-2f6edc0a652b

```
== Related Links

* xref:{relative_path}volume-rebalance-file-move-modify.html[volume
rebalance file-move modify]




[[ID6997b55d3132ff2dfeeaf2b156ef3d85]]
= volume rebalance file-move statistics reset
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset statistics for file-move operations between FlexGroup constituents

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume rebalance file-move statistics reset` command resets completed
statistics for the file-move operations.

The `volume rebalance file-move statistics reset` command is only
supported on FlexGroup volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-volume` parameter are specified, the command
resets statistics for the FlexGroup constituents on the specified volume.
If this parameter is specified by itself, the command resets statistics
for the FlexGroup constituents on all volumes on the specified Vserver.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-vserver` parameter are specified, the command
resets statistics for the FlexGroup constituents on the specified volume.
If this parameter is specified by itself, the command resets statistics
for the FlexGroup constituents on all volumes matching the specified name.

== Examples
```

```
The following example resets completed statistics for all FlexGroup
constituents on the FlexGroup volume "bld_volume" in Vserver "vs0":
```

cluster1::*> volume rebalance file-move statistics show -vserver vs0 -volume bld_volume

Vserver: vs0
Volume: bld_volume
Source Destination
Constituent Files Bytes Files Bytes
--- --- --- --- ---
bld_volume*0001*
In Progress: 1 10GB 0 0B
Completed: 5 100GB 0 0B
*bld_volume*0002
In Progress: 0 0B 1 10GB
Completed: 0 0B 10 150MB
bld_volume*0003*
In Progress: 1 10GB 0 0B
Completed: 15 100GB 0 0B
*bld_volume*0004
In Progress: 0 0B 1 10GB
Completed: 0 0B 10 150MB
bld_volume*0005*
In Progress: 0 0B 1 10GB
Completed: 0 0B 10 150MB
*bld_volume*0006
In Progress: 0 0B 1 10GB
Completed: 0 0B 10 150MB
bld_volume*0007*
In Progress: 0 0B 1 10GB
Completed: 0 0B 10 150MB
*bld_volume*0008
In Progress: 0 0B 1 10GB
Completed: 0 0B 10 150MB

cluster1::*> volume rebalance file-move statistics reset -vserver vs0 -volume bld_volume
(volume rebalance file-move statistics reset)

cluster1::*> volume rebalance file-move statistics show -vserver vs0 -volume bld_volume

Vserver: vs0
Volume: bld_volume
Source Destination
Constituent Files Bytes Files Bytes
--- --- --- --- ---
bld_volume*0001*
In Progress: 1 10GB 0 0B
Completed: 0 0B 0 0B
*bld_volume*0002
In Progress: 0 0B 1 10GB
Completed: 0 0B 0 0B
bld_volume*0003*

In Progress: 1 10GB 0 0B
Completed: 0 0B 0 0B
*bld_volume*0004
In Progress: 0 0B 1 10GB
Completed: 0 0B 0 0B
bld_volume*0005*
In Progress: 0 0B 1 10GB
Completed: 0 0B 0 0B
*bld_volume*0006
In Progress: 0 0B 1 10GB
Completed: 0 0B 0 0B
bld_volume*0007*
In Progress: 0 0B 1 10GB
Completed: 0 0B 0 0B
*bld_volume*0008
In Progress: 0 0B 1 10GB
Completed: 0 0B 0 0B

```
[[IDbdf4e3b1dda1c2fe7674c9e9d9dbe7d5]]
= volume rebalance file-move statistics show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display statistics for file-move operations across FlexGroup Constituents

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume rebalance file-move statistics show` command displays common
statistics for the file-move operations at a FlexGroup constituent level.

The `volume rebalance file-move statistics show` command is only supported
on FlexGroup volumes. For all FlexGroup constituents matching the
specified parameters, the command displays the following information:

* Vserver name
* Volume name
* Constituent name
* Source Files In-Progress: Number of files currently being moved from the
FlexGroup constituent
* Source Files Completed: Number of files already moved from the FlexGroup
constituent
* Source Bytes In-Progress: Total number of bytes of all the files
```

```
currently being moved from the FlexGroup constituent
* Source Bytes Completed: Total number of bytes of all the files already
moved from the FlexGroup constituent
* Destination Files In-Progress: Number of files currently being moved to
the FlexGroup constituent
* Destination Files Completed: Number of files already moved to the
FlexGroup constituent
* Destination Bytes In-Progress: Total number of bytes of all the files
currently being moved to the FlexGroup constituent
* Destination Bytes Completed: Total number of bytes of all the files
already moved to the FlexGroup constituent



You can specify parameters from the above list to display information. For
example, to display statistics only for FlexGroup constituents with more
than one file currently being moved from them, run the command with the `-
source-files-inprogress >1` parameter.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-volume` parameter are specified, the command
displays statistics for the FlexGroup constituents on the specified
volume. If this parameter is specified by itself, the command displays
statistics for the FlexGroup constituents on all volumes on the specified
Vserver.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-vserver` parameter are specified, the command
displays statistics for the FlexGroup constituents on the specified
volume. If this parameter is specified by itself, the command displays
statistics for the FlexGroup constituents on all volumes matching the
specified name.
```

```
``[-constituent <volume name>]`` - Constituent Name [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents matching the specified name.

``[-source-files-inprogress <integer>]`` - Source Files In Progress
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of files currently
being moved from them.

``[-source-files-completed <integer>]`` - Source Files Completed
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of files already
moved from them.

``[-source-bytes-inprogress {<integer>[KB|MB|GB|TB|PB]}]`` - Source Bytes
in Progress [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of bytes currently
being moved from them (across all files currently being moved from them).

``[-source-bytes-completed {<integer>[KB|MB|GB|TB|PB]}]`` - Source Bytes
Completed [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of bytes already
moved from them (across all files already moved from them).

``[-destination-files-inprogress <integer>]`` - Destination Files In
Progress [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of files currently
being moved to them.

``[-destination-files-completed <integer>]`` - Destination Files Completed
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of files already
moved to them.

``[-destination-bytes-inprogress {<integer>[KB|MB|GB|TB|PB]}]`` -
Destination Bytes In Progress [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of bytes currently
being moved to them (across all files currently being moved to them).
```

```
``[-destination-bytes-completed {<integer>[KB|MB|GB|TB|PB]}]`` -
Destination Bytes Completed [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays statistics only for
FlexGroup constituents which have the specified number of bytes already
moved to them (across all files already moved to them).

== Examples

The following example displays statistics for all FlexGroup constituents
on the FlexGroup volume "bld_volume" in Vserver "vs0":
```

cluster1::*> volume rebalance file-move statistics show -vserver vs0 -volume bld_volume

Vserver: vs0
Volume: bld_volume
Source Destination
Constituent Files Bytes Files Bytes
--- --- --- --- ---
bld_volume*0001*
In Progress: 1 10GB 0 0B
Completed: 5 20GB 0 0B
*bld_volume*0002
In Progress: 0 0B 1 10GB
Completed: 0 0B 10 40GB
bld_volume*0003*
In Progress: 1 5GB 0 0B
Completed: 15 100GB 0 0B
*bld_volume*0004
In Progress: 0 0B 1 5GB
Completed: 0 0B 10 80GB
bld_volume*0005*
In Progress: 0 0B 0 0B
Completed: 0 0B 0 0B
*bld_volume*0006
In Progress: 0 0B 0 0B
Completed: 0 0B 0 0B
bld_volume*0007*
In Progress: 0 0B 0 0B
Completed: 0 0B 0 0B
*bld_volume*0008
In Progress: 0 0B 0 0B
Completed: 0 0B 0 0B

```
The following example displays statistics for all FlexGroup constituents
on the FlexGroup volume "bld_volume" in Vserver "vs0" with more than 4
files already sent from them:
```

cluster1::*> volume rebalance file-move statistics show -vserver vs0 -volume bld_volume -source-files -completed >4

Vserver: vs0
Volume: bld_volume
Source Destination
Constituent Files Bytes Files Bytes
--- --- --- --- ---
bld_volume*0001*
In Progress: 1 10GB 0 0B
Completed: 5 20GB 0 0B
*bld_volume*0003
In Progress: 1 5GB 0 0B
Completed: 15 100GB 0 0B

```
:leveloffset: -1


= volume recovery-queue commands

:leveloffset: +1


[[ID4b40a5662e40e892b64048b0a27afcd7]]
= volume recovery-queue modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify attributes of volumes in the recovery queue

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume recovery-queue modify` command modifies the retention period
of a volume in the recovery queue.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume that is to be modified.
```

```
``[-retention-hours <integer>]`` - Volume Delete Retention [.advaced-
priv]#(privilege: advanced)#::
This specifies the time in hours that the volume will be available in the
recovery queue for recovery.For example, a value of "10" means that the
volume will be retained in the recovery queue for at least 10 hours. The
maximum allowed Volume retention-hours is 4294967295 hours.

== Examples

The following example modifies a volume named vol1_1234 on a Vserver named
vs1 in the recovery queue.
```

cluster1::*> volume recovery-queue modify -vserver vs1 -volume vol1_1234 -retention-hours 10

```
[[ID794755ba7337cd572b2fae9d57664bd8]]
= volume recovery-queue purge-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Purge all volumes from the recovery queue belonging to a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume recovery-queue purge-all` command purges all volumes from the
recovery queue belonging to a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver whose volumes in the recovery queue need to be
purged.

== Examples

The following example purges all volumes on a Vserver named vs1 from the
recovery queue.
```

cluster1::*> volume recovery-queue purge-all -vserver vs1

```
[[ID0f171c50262e9c11b52e0c31b7dee8ac]]
= volume recovery-queue purge
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Purge volumes from the recovery queue belonging to a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume recovery-queue purge` command purges a volume from the
recovery queue belonging to a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume that is to be purged from the recovery queue.

== Examples

The following example purges a volume named vol1_1234 on a Vserver named
vs1 from the recovery queue.
```

cluster1::*> volume recovery-queue purge -vserver vs1 -volume vol1_1234

```
[[ID24b8b25fe3fc99bc7bd4d31cbc9cec40]]
= volume recovery-queue recover-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Recover all volumes from the recovery queue belonging to a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume recovery-queue recover-all` command recovers all the volumes
from the recovery-queue belonging to a Vserver.

NOTE: This command will restore volumes with their previous access types
and volumes will remain offline. This command will not restore the volumes
with their previous name, junction, LUN mappings, NVMe namespace mappings,
quota policy rules and snapshot schedules.



== Parameters

``-vserver <vserver name>`` - Vserver name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver for which the volumes need to be recovered.

== Examples

The following example recovers all volumes on a Vserver named vs1 from the
recovery queue.
```

cluster1::*> volume recovery-queue recover-all -vserver vs1

```
[[ID8719732c3f753d08729a45478edcdbf6]]
= volume recovery-queue recover
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Recover volumes from the recovery queue belonging to a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume recovery-queue recover` command recovers a volume from the
recovery queue belonging to a Vserver.

NOTE: This command will restore the volume with its previous access type
and the volume will remain offline. This command will not restore the
volume with its previous name, junction, LUN mappings, NVMe namespace
mappings, quota policy rules and snapshot schedules.



== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the volume that is to be recovered.

== Examples

The following example recovers a volume named vol1_1234 on a Vserver named
vs1 from the recovery queue.
```

cluster1::*> volume recovery-queue recover -vserver vs1 -volume vol1_1234

```
[[ID7a09cfbe886030e1d7e9d0b5c4ed2bae]]
= volume recovery-queue show
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show volumes in the recovery queue

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `volume recovery-queue show` command displays information about
volumes in the recovery queue. The command output depends on the
parameter(s) specified with the command. If no parameters are specified,
the command displays the following information about all volumes:

* Vserver name
* Volume name
* Deletion Request Time
* Retention Hours



To display information about a single volume, run the command with the `-
vserver` and `-volume` parameters.

To display detailed information about all volumes, run the command with
the `-instance` parameter.

You can specify additional parameters to display information that matches
only those parameters. For example, to display information only about
volumes with retention hours 10, run the command with the `-retention-
hours 10` parameter.

== Parameters

``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed. The fields Vserver
and Volume are the default fields.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is specified, the command displays detailed information
about all fields.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
```

```
If this parameter and the `-volume` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter and the `-vserver` parameter are specified, the command
displays detailed information about the specified volume. If this
parameter is specified by itself, the command displays information about
all volumes matching the specified name.

``[-delete-time <Date>]`` - Deletion Request Time [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified delete time.

``[-retention-hours <integer>]`` - Volume Delete Retention [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified retention hours.

``[-instance-uuid <UUID>]`` - Volume Instance UUID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the volume or volumes that match the specified instance uuid.

== Examples

The following example displays information about all volumes on a Vserver
named vs1 in the recovery queue:
```

cluster1::*> volume recovery-queue show -vserver vs1
Vserver Volume Deletion Request Time Retention Hours
--------- ----------- -------------------- ---------------
vs1 vol1_1234 Wed Aug 22 06:50:19 2012 10
vs1 vol2_1235 Wed Aug 22 06:50:26 2012 12
2 entries were displayed.

```
The following example displays detailed information about a volume named
vol1_1234 on a Vserver named vs1 in the recovery queue:
```

cluster1::*> volume recovery-queue show -vserver vs1 -volume vol1_1234
Vserver Name: vs1
Volume Name: vol1_1234
Deletion Request Time: Wed Aug 22 06:50:19 2012

Volume Delete Retention: 10

```
:leveloffset: -1


= volume schedule-style commands

:leveloffset: +1


[[IDf45a8379b53a8da002fe5d4cd632ede3]]
= volume schedule-style prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disables volume schedule style feature and sets schedule style to default
(create-time)

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command will disable the volume schedule style feature and set
schedule style to default (create-time).

== Examples

The following example prepares the schedule-style on all volumes for
revert/downgrade.
```

cluster1::*> volume schedule-style prepare-to-downgrade

```
:leveloffset: -1


= volume snaplock commands

:leveloffset: +1


[[IDfe91f8f6edff812f4e47c7d20c473822]]
= volume snaplock modify
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify SnapLock attributes of a SnapLock volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 The ` volume snaplock modify ` command modifies one or more SnapLock
attributes of a SnapLock volume.

== Parameters

``-vserver <vserver name>`` - Vserver::
 This specifies the vserver which owns the required SnapLock volume.

``-volume <volume name>`` - Volume::
 This specifies the SnapLock volume whose attribute needs to be modified.

``[-minimum-retention-period {{<integer>
seconds|minutes|hours|days|months|years} | infinite}]`` - Minimum
Retention Period::
Specifies the minimum allowed retention period for files committed to WORM
state on the volume. Any files committed with a retention period shorter
than this minimum value, is assigned this minimum value.
+
If this option value is ``_infinite_`` , then every file committed to the
volume will have a retention period that never expires.
+
Otherwise, the retention period is specified as a number followed by a
suffix. The valid suffixes are seconds, minutes, hours, days, months, and
years. For example, a value of 6months represents a retention period of 6
months. The maximum allowed retention period is 70 years. This option is
not applicable while extending retention period of an already committed
WORM file.

``[-default-retention-period {{<integer>
seconds|minutes|hours|days|months|years} | min | max | infinite |
unspecified}]`` - Default Retention Period::
Specifies the default retention period that is applied to files while
committing to WORM state without an associated retention period.
```

```
+
If this option value is ``_min_`` , then minimum-retention-period is used
as the default retention period. If this option value is ``_max_`` , then
maximum-retention-period is used as the default retention period. If this
option value is ``_infinite_`` , then a retention period that never
expires will be used as the default retention period. If this option value
is ``_unspecified_`` , then the file will be retained forever; however,
the retention time can be changed to an absolute value.
+
The retention period can also be explicitly specified as a number followed
by a suffix. The valid suffixes are seconds, minutes, hours, days, months,
and years. For example, a value of 6months represents a retention period
of 6 months. The maximum valid retention period is 70 years.

``[-maximum-retention-period {{<integer>
seconds|minutes|hours|days|months|years} | infinite}]`` - Maximum
Retention Period::
Specifies the maximum allowed retention period for files committed to WORM
state on the volume. Any files committed with a retention period longer
than this maximum value, is assigned this maximum value.
+
If this option value is ``_infinite_`` , then files that have retention
period that never expires might be committed to the volume.
+
Otherwise, the retention period is specified as a number followed by a
suffix. The valid suffixes are seconds, minutes, hours, days, months, and
years. For example, a value of 6months represents a retention period of 6
months. The maximum allowed retention period is 70 years. This option is
not applicable while extending retention period of an already committed
WORM file.

``[-autocommit-period {{<integer> minutes|hours|days|months|years} |
none}]`` - Autocommit Period::
Specifies the autocommit period for SnapLock volume. All files which are
not modified for a period greater the autocommit period of the volume are
committed to WORM state.
+
The autocommit period option is specified as a number followed by a
suffix. The valid suffixes for autocommit period are hours, minutes, days,
months and years. For example, a value of 2hours represents an autocommit
period of 2 hours. The minimum allowed autocommit period is 5 minutes and
the maximum allowed autocommit period is 10 years.
+
If this option value is ``_none_`` , then autocommit is disabled on the
SnapLock volume.
```

```
``[-is-volume-append-mode-enabled {true|false}]`` - Is Volume Append Mode
Enabled::
Specifies if the volume append mode is enabled or disabled.
+
It can be modified only when the volume is not mounted and does not have
any data or snapshots.
+
The volume append mode is not supported on SnapLock audit log volumes.
+
When it is enabled, all the files created with write permissions on the
volume are WORM appendable files by default. All the WORM appendable files
that are not modified for a period greater than the autocommit period of
the volume are also committed to the WORM read-only state.
+
If it is set to ``_true_`` , then the volume append mode is enabled.
+
If it is set to ``_false_`` , then the volume append mode is disabled.
+
Volume append mode is disabled by default when the volume is created.

``[-privileged-delete {disabled|enabled|permanently-disabled}]`` -
Privileged Delete::
Specifies the privileged-delete attribute of a SnapLock volume. This
parameter must be specified alone.
+
If it is set to ``_enabled_`` then the privileged-delete operation can be
performed using the xref:{relative_path}volume-file-privileged-
delete.html[ volume file privileged-delete ] command.
+
If it is set to ``_disabled_`` , then the privileged-delete operation is
not supported.
+
Once it is set to ``_permanently-disabled_`` , then neither the
privileged-delete operation nor any change in the volume privileged-delete
attribute is permitted.

== Examples

The following command sets `-default-retention-period` of a given SnapLock
volume:
```

cluster1::> volume snaplock modify -volume vol_slc -default-retention-period 2years

cluster1::>

```
The following command sets `-maximum-retention-period` of a given SnapLock
volume to ``_infinite_`` :
```

cluster1::> volume snaplock modify -volume vol_slc -maximum-retention-period infinite

cluster1::>

```
The following command enables the privileged-delete operation on a
SnapLock volume.
```

cluster1::> volume snaplock modify -vserver vs1 -volume vol_sle -privileged-delete enabled
[Job 38] Job succeeded: Privileged-delete Attribute Change for Volume "vs1:vol_sle" Completed.
cluster1::>
cluster1::>volume snaplock show -vserver vs1 -volume vol_sle -fields privileged-delete
vserver volume privileged-delete
------- -------- -----------------
vs1 vol_sle enabled

cluster1::>

```
== Related Links

* xref:{relative_path}volume-file-privileged-delete.html[volume file
privileged-delete]




[[ID9afbb242421a47f732fb86a4e42e891b]]
= volume snaplock prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Prepares the system for downgrade

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume snaplock prepare-to-downgrade` command prepares nodes to
downgrade to a release without SnapLock volume append mode feature. Prior
to disabling the feature, the command disables volume append mode on all
SnapLock volumes in the cluster.

== Examples

The following example disables the SnapLock volume append mode feature in
the local cluster:
```

cluster1::> volume snaplock prepare-to-downgrade

```
[[ID0d2cd1bb54930aeffeaa4cf2b910248e]]
= volume snaplock show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display SnapLock attributes of a SnapLock volume
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 The ` volume snaplock show ` command displays following information :

* Vserver name
* Volume name
* SnapLock Type of the volume
* Minimum retention period applicable of the volume
* Default retention period applicable of the volume
* Maximum retention period applicable of the volume
* Autocommit period of the volume
* Volume Append Mode attribute of the volume
* Privileged Delete attribute of the volume
* Litigation count on the volume
* Volume expiry time of the volume
* Volume ComplianceClock
* SnapLock audit log volume
* Unspecified retention file count on the volume

This command is applicable only for SnapLock volumes.

== 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.

``[-vserver <vserver name>]`` - Vserver::
 If this parameter is specified, the command displays information for all
the SnapLock volumes that match the specified `-vserver` value.

``[-volume <volume name>]`` - Volume::
 If this parameter is specified, the command displays information for the
specified `-volume` value.

``[-type {non-snaplock|compliance|enterprise}]`` - SnapLock Type::
 If this parameter is specified, the command displays all the volumes that
match the specified `-type` value.
```

```
``[-minimum-retention-period {{<integer>
seconds|minutes|hours|days|months|years} | infinite}]`` - Minimum
Retention Period::
If this parameter is specified, the command displays all the volumes that
match the specified `-minimum-retention-period` value.

``[-default-retention-period {{<integer>
seconds|minutes|hours|days|months|years} | min | max | infinite |
unspecified}]`` - Default Retention Period::
If this parameter is specified, the command displays all the volumes that
match the specified `-default-retention-period` value.

``[-maximum-retention-period {{<integer>
seconds|minutes|hours|days|months|years} | infinite}]`` - Maximum
Retention Period::
If this parameter is specified, the command displays all the volumes that
match the specified `-maximum-retention-period` value.

``[-autocommit-period {{<integer> minutes|hours|days|months|years} |
none}]`` - Autocommit Period::
If this parameter is specified, the command displays all the volumes that
match the specified `-autocommit-period` value.

``[-is-volume-append-mode-enabled {true|false}]`` - Is Volume Append Mode
Enabled::
If this parameter is specified, the command displays all the volumes that
match the specified `-is-volume-append-mode-enabled` value.

``[-privileged-delete {disabled|enabled|permanently-disabled}]`` -
Privileged Delete::
If this parameter is specified, the command displays all the volumes that
match the specified `-privileged-delete` value.

``[-expiry-time <text>]`` - Expiry Time::
 If this parameter is specified, the command displays all the volumes that
match the specified `-expiry-time` value.

``[-compliance-clock-time <text>]`` - ComplianceClock Time::
 If this parameter is specified, the command displays all the volumes that
match the specified `-compliance-clock-time` value.

``[-litigation-count <integer>]`` - Litigation Count::
If this parameter is specified, the command displays all the volumes that
match the specified `-litigation-count` value.

``[-is-audit-log-volume {true|false}]`` - Is SnapLock Audit Log Volume::
If this parameter is specified, the command displays all the volumes that
```

```
match the specified `-is-audit-log-volume` value.

``[-unspecified-retention-file-count <integer>]`` - Unspecified Retention
File Count::
If this parameter is specified, the command displays all the volumes that
match the specified `-unspecified-retention-file-count` value.

== Examples

The following command shows summary of SnapLock volumes on a vserver:
```

cluster1::> volume snaplock show
Vserver Volume SnapLock Type ComplianceClock Time
---------- --------- ---------- ----------
vs1 vol_slc compliance Mon Jan 19 14:12:34 IST 2015 +05:30
vs1 vol_sle enterprise Mon Jan 19 14:12:34 IST 2015 +05:30
2 entries were displayed.

cluster1::>

```
The following commands lists the complete SnapLock attributes of two given
SnapLock volumes:
```

cluster1::> volume snaplock show -vserver vs1 -volume vol_slc
Vserver Name: vs1
Volume Name: vol_slc
SnapLock Type: compliance
Minimum Retention Period: 1 years
Default Retention Period: max
Maximum Retention Period: 30 years
Autocommit Period: 12 hours
Is Volume Append Mode Enabled: false
Privileged Delete: permanently-disabled
Expiry Time: Thu May 11 14:37:21 GMT 2017
ComplianceClock Time: Wed May 11 20:08:41 IST 2016 +05:30
Litigation Count: 0
Is SnapLock Audit Log Volume: false
Unspecified Retention File Count: 0

cluster1::>

cluster1::> volume snaplock show -vserver vs1 -volume vol_sle
Vserver Name: vs1
Volume Name: vol_sle
SnapLock Type: enterprise
Minimum Retention Period: 6 months
Default Retention Period: min
Maximum Retention Period: infinite
Autocommit Period: none
Is Volume Append Mode Enabled: false

Privileged Delete: enabled
Expiry Time: infinite
ComplianceClock Time: Wed May 11 20:08:44 IST 2016 +05:30
Litigation Count: 0
Is SnapLock Audit Log Volume: false
Unspecified Retention File Count: 0

```
:leveloffset: -1


= volume snapshot commands

:leveloffset: +1


[[ID1ec76f07661b2d5fc7f1a8ec549e8d7f]]
= volume snapshot compute-reclaimable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Calculate the reclaimable space if specified snapshots are deleted

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot compute-reclaimable` command calculates the volume
space that can be reclaimed if one or more specified snapshots are
deleted.

The command heavily uses system’s computational resources so it can reduce
the performance for client requests and other system processes. Therefore,
the queries that use queries that use query operators (`*` , `|` , etc.),
are disabled for this command. You should not specify more than three
snapshots per query. snapshots must be specified as a comma-separated list
with no spaces after the commas.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
```

```
This specifies the volume for which reclaimable space is to be calculated.

``-snapshots <snapshot name>,...`` - List of Snapshots::
This specifies one or more than one snapshots that are to be considered
for deletion. If you list more than one snapshot, specify a comma-
separated list with no spaces after the commas.

== Examples

The following example calculates the space that can be reclaimed if the
snapshot named hourly.2008-01-10_1505 is deleted on a volume named vol3,
which is a part of the Vserver named vs0:
```

cluster1::> volume snapshot compute-reclaimable -vserver vs0
-volume vol3 -snapshots hourly.2008-01-10_1505

```
[[ID4ed2f6e60b191b670e99561565df24fd]]
= volume snapshot create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a snapshot

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot create` command creates a snapshot of a specified
volume.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver that contains the volume on which the snapshot
is to be created.

``-volume <volume name>`` - Volume::
This specifies the volume where a snapshot is to be created.

``-snapshot <snapshot name>`` - Snapshot::
This specifies the name of the snapshot that is to be created.
```

```
``[-comment <text>]`` - Comment::
This optionally specifies a comment for the snapshot.

``[-foreground {true|false}]`` - Foreground Process::
If you use this option and select ``_false_`` , the snapshot creation
process runs in the background. If you use this option and select
``_true_`` , the snapshot creation process runs in the foreground. The
default is ``_true_`` .

``[-snapmirror-label <text>]`` - Label for SnapMirror Operations::
If you specify this option, the snapshot is created with the SnapMirror
Label that you specify. If this option is not specified, the snapshot is
created with no SnapMirror Label. The SnapMirror Label is used by the
Vaulting subsystem when you back up snapshots to the Vault Destination.

``[-expiry-time <MM/DD/YYYY HH:MM:SS>]`` - Expiry Time::
If you specify this option, the snapshot is created with the expiry time
that you specify. The expiry time indicates the time at which the snapshot
becomes eligible for deletion.

``[-snaplock-expiry-time {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm] | infinite}]``
- SnapLock Expiry Time::
This optionally specifies an expiry time for the snapshot. A snapshot with
an expiry time cannot be deleted until the expiry time has elapsed. This
option takes effect only on volumes with ``_snapshot-locking-enabled_``
set to true. If this option is not specified, the snapshot will not be
under retention.

== Examples

The following example creates a snapshot named vol3_snap on a volume named
vol3 on a Vserver named vs0. The snapshot is given the comment "Single
snapshot" and the operation runs in the background.
```

cluster1::> volume snapshot create -vserver vs0 -volume vol3 -snapshot vol3_snapshot -comment "Single snapshot" -foreground false

```
[[ID073384bbe29b6445d2a8c3dbb431faaa]]
= volume snapshot delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a snapshot
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot delete` command deletes a snapshot from a specified
volume.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver that contains the volume on which the specified
snapshot is saved.

``-volume <volume name>`` - Volume::
This specifies the volume from which a snapshot is to be deleted.

``-snapshot <snapshot name>`` - Snapshot::
This specifies the snapshot that is to be deleted.

``[-foreground {true|false}]`` - Foreground Process::
If you use this option and set it to ``_false_`` , the delete operation
runs as a background process. If you specify this option and set it to
``_true_`` , the operation runs as a foreground process. The default is
``_true_`` .

``[-force <true>]`` - Force Delete [.advaced-priv]#(privilege:
advanced)#::
If you use this switch, the snapshot is immediately deleted without
generating any confirmation messages. If you do not use this option the
operation generates confirmation messages and the operation is disallowed
on application tagged volumes. Passing in a value of ``_true_`` is
supported, but not required. The `force` switch is typically used for
scripting applications where users cannot directly confirm the delete
operation.

``[-ignore-owners <true>]`` - Ignore Snapshot Owners [.advaced-
priv]#(privilege: advanced)#::
If you use this switch, the command ignores other processes that might be
accessing the snapshot. If you do not use this option the operation
exhibits default behavior and checks the owners tags before allowing the
deletion to occur. Passing in a value of ``_true_`` is supported, but not
required.

== Examples
```

```
The following example deletes a snapshot named vol3_daily from a volume
named vol3 on a Vserver named vs0:
```

cluster1::> volume snapshot delete -vserver vs0 -volume vol3 -snapshot vol3_daily

```
[[ID80f4bccd8db84f8828d526324ea6a33a]]
= volume snapshot modify-snaplock-expiry-time
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify expiry time of a SnapLock Snapshot copy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot modify-snapshot-expiry-time` extends SnapLock expiry
time of an existing snapshot.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver that contains the volume on which the snapshot
is located.

``-volume <volume name>`` - Volume::
This specifies the volume where a snapshot is to be located.

``-snapshot <text>`` - Snapshot::
This specifies the name of the snapshot locked by SnapLock whose SnapLock
expiry time needs to be modified.

``[-expiry-time {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm] | infinite}]`` -
SnapLock Expiry Time::
Specifies the new SnapLock expiry that is applied to snapshot locked by
SnapLock.
+
If this option value is ``_infinite_`` , then a retention period that
never expires is applied to the snapshot.

== Examples

The following example extends the retention period of a snapshot
``_snap1_`` to ``_"03/03/2020 00:00:00"_`` :
```

cluster1::> volume snapshot modify-snaplock-expiry-time -vserver vs1 -volume vol1 -snapshot snap1 -expiry

-time "03/03/2020 00:00:00"
cluster1::>

```
The following example extends the retention period of a snapshot
``_snap2_`` to ``_infinite_`` :
```

cluster1::> volume snapshot modify-snaplock-expiry-time -vserver vs1 -volume vol1 -snapshot snap2 -expiry -time infinite
cluster1::>

```

```

cluster1::> volume snapshot show -vserver vs1 -fields snaplock-expiry-time
vserver volume snapshot snaplock-expiry-time
------- ------ -------- ------------------------
vs1 vol1 snap1 3/3/2020 00:00:00 +05:30
vs1 vol1 snap2 infinite

```
[[IDe4d423e537deb9f5834a85bfbcd5a58d]]
= volume snapshot modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify snapshot attributes

*Availability:* This command is available to _cluster_ and _Vserver_ administrators at the _admin_ privilege level.

== Description

The `volume snapshot modify` command enables you to change the text comment associated with a snapshot.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver that contains the volume on which the specified snapshot is saved.

``-volume <volume name>`` - Volume::
This specifies the volume whose snapshot is to be modified.

``-snapshot <snapshot name>`` - Snapshot::
```

```
This specifies the snapshot whose text comment is to be modified.

``[-comment <text>]`` - Comment::
This specifies the new comment for the snapshot.

``[-snapmirror-label <text>]`` - Label for SnapMirror Operations::
This specifies the SnapMirror Label for the snapshot. The SnapMirror Label
is used by the Vaulting subsystem when you back up snapshots to the Vault
Destination. If an empty label ("") is specified, the existing label will
be deleted.

``[-expiry-time <MM/DD/YYYY HH:MM:SS>]`` - Expiry Time::
This specifies the expiry time for the snapshot. The expiry time indicates
the time at which the snapshot becomes eligible for deletion. If an expiry
time of ("0") is specified, the existing expiry time will be deleted.

== Examples

The following example modifies the comment of a snapshot named
vol3_snapshot of a volume named vol3 on a Vserver named vs0. The comment
is changed to "Pre-upgrade snapshot".
```

cluster1::> volume snapshot modify -vserver vs0 -volume vol3
-snapshot vol3_snapshot -comment "Pre-upgrade snapshot"

```
[[IDa2d18c57806b72f4d66df1ed0615f8f0]]
= volume snapshot partial-restore-file
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore part of a file from a snapshot

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot partial-restore-file` command enables you to restore
a range of bytes in a file from the version of the file saved in the
snapshot. This command is intended to be used to restore particular pieces
of LUNs, NVMe namespaces, and NFS or CIFS container files that are used by
a host to store multiple sources of data. For example, a host might be
storing multiple user databases in the same LUN. A partial file restore
```

```
can be used to restore one of those databases in the LUN without touching
other databases stored in the LUN. This command is not intended for
restoring parts of normal user-level files that are stored in the volume.
You should use xref:{relative_path}volume-snapshot-restore-
file.html[volume snapshot restore-file] command to restore normal user-
level files. The volume for the partial-restore should be online during
this operation.

For LUNs and NVMe namespaces, this command is supported across all LUN and
NVMe namespace source and destination objects with equal logical block
sizes.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver which contains the volume.

``[-volume <volume name>]`` - Volume Name::
This specifies the volume in which the snapshot is saved.

``-s, -snapshot <snapshot name>`` - Snapshot Name::
This specifies the snapshot which contains the version of file from which
a range of bytes is restored. The source file, LUN, or NVMe namespace must
be present in the snapshot.

``-path <text>`` - Filepath::
This specifies the relative path to the file, LUN, or NVMe namespace which
is partially restored from the snapshot. You should specify the `-volume`
option so that the file, LUN, or NVMe namespace is searched and restored
from the snapshot of the specified volume. If you do not specify the `-
volume` then the file, LUN, or NVMe namespace is searched and restored
from the snapshot of the root volume. The destination file, LUN, or NVMe
namespace must be present in the active file system.

``-start-byte <integer>`` - Starting Byte Offset (Multiple of 4096)::
This specifies the starting byte offset in the file to partially restore.
The first byte of the file is byte zero. The start byte must be a multiple
of 4096. In addition, the start byte must not exceed the size of the
source or destination file.

``-byte-count <integer>`` - Number of Bytes to Restore (Multiple of
4096)::
This specifies the total number of bytes to restore, beginning at the `-
start-byte` value. The `-byte-count` option must be a multiple of 4096.
The maximum number of bytes that can be restored is 16 MB. The byte count
must not exceed the range of the source or destination file.
```

```
== Examples

The following example restores first 4096 bytes in the file ``_foo.txt_``
inside the volume ``_vol3_`` from the snapshot ``_vol3_snap_`` :
```

cluster1::> volume snapshot partial-restore-file -vserver vs0 -volume vol3 -snapshot vol3_snap -path /foo.txt -start-byte 0 -byte-count 4096

```
== Related Links

* xref:{relative_path}volume-snapshot-restore-file.html[volume snapshot
restore-file]




[[ID7edf5bde1128bfa52c35722fdf1a0669]]
= volume snapshot prepare-for-revert
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Deletes multiple Snapshot copies of the current File System version.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command will delete all snapshots that have the format used by the
current version of ONTAP. It will fail if snapshot polices are enabled, or
if snapshots have an owner.

NOTE: Snapshot policies must be disabled prior to running this command.



== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
The name of the node.

== Examples

The following example prepares the snapshots for revert.
```

cluster1::*> volume snapshot prepare-for-revert -node node1

```
[[ID617d10698898a4abad436282a5724e1e]]
= volume snapshot rename
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a snapshot

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot rename` command renames a snapshot.

NOTE: You cannot rename a snapshot that is created as a reference copy
during the execution of the `volume move` command.



== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver that contains the volume on which the specified
snapshot is to be renamed

``-volume <volume name>`` - Volume::
This specifies the volume that contains the snapshot to be renamed.

``-snapshot <snapshot name>`` - Snapshot::
This specifies the snapshot that is to be renamed.

``-new-name <snapshot name>`` - Snapshot New Name::
This specifies the new name for the snapshot.

``[-force <true>]`` - Force Rename [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the snapshot rename operation is allowed
on application tagged volumes. Otherwise, the operation is disallowed on
application tagged volumes.

== Examples

The following example renames a snapshot named vol3_snap on a volume named
vol3 and a Vserver named vs0. The snapshot is renamed to
vol3_snap_archive.
```

cluster1::> volume snapshot rename -vserver vs0 -volume vol3 -snapshot vol3_snap -new-name vol3_snap_archive

```
[[ID67db90216456251c3c6466ac79d27524]]
= volume snapshot restore-file
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore a file from a snapshot

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot restore-file` command enables you to restore a single
file to a version saved in the snapshot. You can restore a file over an
existing copy of the file in the parent read-write volume or to a
different location within the same parent read-write volume. If the
destination file for the restore operation does not exist, a new file is
created with the same version as the one saved in the snapshot. If the
destination file for the restore operation exists, then it is overwritten
by the version from the snapshot. This operation is used to restore normal
user-level files, LUNs and NVMe namespaces. The command also supports
restoring normal user-level files with streams. The command fails if you
try to restore directories (and their contents). During the restore
operation the parent read-write volume should remain online. The command
fails if the destination path for the restore operation is in a different
volume than the source volume.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver which contains the volume.

``[-volume <volume name>]`` - Volume Name::
This specifies the volume which contains the specified snapshot.

``-s, -snapshot <snapshot name>`` - Snapshot Name::
This specifies the snapshot from which the file is restored.

``-path <text>`` - Filepath::
This specifies the relative path to the file which is restored from the
```

```
snapshot. You should specify the `-volume` option so that the file is
searched and restored from the snapshot of the specified volume. If you do
not specify the `-volume` then the file is searched and restored from the
snapshot of the root volume.

``[-r, -restore-path <text>]`` - Restore Filepath::
This option specifies the destination location inside the volume where the
file is restored. If you do not specify this option, the file is restored
at the same location referred by `-path` option. If you specify `-restore-
path` option, then it should refer to a relative path location within the
same volume which contains the source file. If you do not specify `-
volume` along with the relative path, the file is restored in the root
volume.

``[-split-disabled <true>]`` - Disable Space Efficient LUN Splitting::
If you use this option and set it to ``_true_`` , space efficient LUN or
NVMe namespace clone split is not allowed during the restore operation. If
you use this option and set it to ``_false_`` or do not use this option,
then space efficient LUN or NVMe namespace clone split is allowed during
the restore operation.

``[-ignore-streams <true>]`` - Ignore Streams::
If you use this parameter, the file is restored without its streams. By
default, the streams are restored.

== Examples

The following example restores a file ``_foo.txt_`` from the snapshot
``_vol3_snap_`` inside the volume ``_vol3_`` contained in a Vserver
``_vs0_`` :
```

cluster1::> volume snapshot restore-file -vserver vs0 -volume vol3 -snapshot vol3_snap -path /foo.txt

```
[[ID1a73add35d612f0c7d21949444eec3ec]]
= volume snapshot restore
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore the volume to a snapshot.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `volume snapshot restore` command restores a snapshot to be the read-
write parent volume for the volume family. This replaces the current
working copy of the volume with the snapshot that results in a loss of all
changes made since the snapshot was created.

NOTE: You should manually update all the SnapMirror relationships of a
volume immediately after you restore its snapshot. Not doing so can result
in unusable SnapMirror relationships that must be deleted and re-created.

After the restore is complete, the size of the flexible volume will be set
to either the current volume size or the snapshot size - whichever is
greater.



== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver that contains the volume on which the specified
snapshot to be restored is saved.

``-volume <volume name>`` - Volume::
This specifies the parent read-write volume whose snapshot is to be
restored to take its place.

``-snapshot <snapshot name>`` - Snapshot::
This specifies the snapshot that is to be restored to be the read-write
parent volume.

``[-force <true>]`` - Force Restore::
If you use this parameter, the snapshot is restored even if the volume has
one or more newer snapshots which are currently used as reference snapshot
by SnapMirror. If a restore is done in this situation, this will cause
future SnapMirror transfers to fail. The SnapMirror relationship may be
repaired using xref:{relative_path}snapmirror-resync.html[snapmirror
resync] command if a common snapshot is found between the source and
destination volume. If there is no common snapshot between the source and
the destination volume, a baseline SnapMirror copy would be required. If
you use this parameter, the operation is also allowed on application
tagged volumes.

``[-preserve-lun-ids {true|false}]`` - Preserve LUN Identifiers::
This option enables you to select whether the snapshot restore needs to be
non-disruptive to clients due to LUN or NVMe namespace identifiers
changing. If you use this option and set it to ``_true_`` , or choose to
```

```
not use this option at all, the `volume snapshot restore` command fails if
the system determines that it cannot be non-disruptive with regards to LUN
or NVMe namespace identifiers. If you use this option and set it to
``_false_`` , the restore operation proceeds even if this might cause
client-visible effects. In this case, administrators should take the LUNs
or NVMe namespaces offline before proceeding.

== Examples

The following example restores a snapshot named vol3_snap_archive to be
the parent read-write volume for the volume family. The existing read-
write volume is named vol3 and is located on a Vserver named vs0:
```

cluster1::> volume snapshot restore -vserver vs0 -volume vol3
-snapshot vol3_snap_archive

```
== Related Links

* xref:{relative_path}snapmirror-resync.html[snapmirror resync]




[[IDaaafae59478d75f63784c3d6d9a68a7f]]
= volume snapshot show-delta
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Computes delta between two Snapshot copies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot show-delta` command returns the number of bytes that
changed between two snapshots or a snapshot and the active filesystem.
This is calculated from the number of blocks that differ multiplied by the
block size. The command also shows the time elapsed between the snapshots
in seconds.

Queries that use query operators (`*` , `|` , etc.) are disabled for this
command to avoid performance degradation for client requests.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the volume for which the delta is to be calculated.

``-snapshot1 <snapshot name>`` - First Snapshot Name::
This specifies the first snapshot for the comparison.

``[-snapshot2 <snapshot name>]`` - Second Snapshot Name::
This specifies the second snapshot for the comparison. If the field is not
specified, it is assumed to be the Active File System.

== Examples

The following example shows the bytes changed and the time separating the
two snapshots:
```

cluster1::> volume snapshot show-delta -vserver vs0 -volume vol2 -snapshot1 one -snapshot2 two
A total of 139264 bytes (34 blocks) are different. Elapsed time between the snapshots: 1s.

```
[[ID123983f8f247b395d40614f17279e055]]
= volume snapshot show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of snapshots

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot show` command displays information about snapshots.
The command output depends on the parameters specified with the command.
If no parameters are specified, the command displays a table with the
following information about all the available snapshots:

* Vserver name
* Volume name
```

```
* Snapshot name
* State
* Size
* Percentage of total blocks in the parent volume
* Percentage of used blocks in the parent volume

To display a detailed list view with additional information, run the
command and select the `-instance` view. In addition to the above
mentioned information about the snapshots, the detailed list view provides
the following additional information:

* Creation time
* Snapshot busy
* List of the snapshot's owners
* Comment associated with the snapshot
* SnapMirror Label associated with the snapshot
* 7-Mode Snapshot
* Constituent Snapshot
* Expiry Time
* SnapLock Expiry Time
* SnapLock Snapshot Expired
* Seconds until SnapLock Snapshot Expiry

At the advanced or higher privilege level the detailed view provides the
following additional information:

* Snapshot's Dataset ID
* Snapshot's master Dataset ID
* Number of consistency points in the snapshot
* Internal status of the snapshot
* File system version
* File system block format
* Physical Snap ID
* Logical Snap ID
* Database record owner
* Snapshot tags
* Instance UUID
* Version UUID
* Node
* AFS used size
* Compression savings size
* Deduplication savings size
* Vbn0 savings size
* Reserved size
* Logical used size
* Performance metadata size
```

```
* Status of FlexGroup Qtree support in the snapshot



NOTE: For snapshots whose parent volume is a FlexGroup, some information
is not available and empty values will be displayed. This information
includes:

* State
* Size
* Percentage of total blocks in the parent volume
* Percentage of used blocks in the parent volume

 All information is available for snapshots whose parent volume is a
FlexGroup Constituent.

At the admin and advanced privilege level, snapshots whose parent volume
is a FlexGroup Constituent are not displayed by default. To display these,
run the command and set the `is-constituent` to ``_true_`` . At the
diagnostic or higher privilege level, all snapshots are displayed by
default.

The list view is automatically enabled if a single snapshot is specified
by using the `-vserver` , `-volume` and `-snapshot` options together.

A preformatted query for displaying the time-related information is
available by specifying the `-time` format specifier. This displays a
table that contains the following fields for all the available snapshots:

* Vserver name
* Volume name
* Snapshot name
* Creation time



By using the `-fields` option you can choose to print only the certain
fields in the output. This presents the selected fields in a table view.
This is ideal when you want additional information to be different from
the information that is provided by the default table view, but would like
it in a format which is visually easy to compare.

You can specify additional parameters to display the information that
matches only those parameters. For example, to display information only
about snapshots of the load-sharing volumes, run the command with the `-
volume-type LS` parameter. If you specify multiple filtering parameters,
only those snapshots that match all the specified parameters are
```

```
displayed.

== 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.

| ``[-time  ]``::
If the `-time` format is specified, the command displays time related
information about all entries.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver::
If you use this parameter, the snapshots located only on the specified
Vserver will be displayed.

``[-volume <volume name>]`` - Volume::
If you use this parameter only snapshots located on the specified volume
will be displayed.

``[-snapshot <snapshot name>]`` - Snapshot::
If you use this parameter only snapshots matching the specified name will
be displayed.

``[-dsid <integer>]`` - Snapshot Data Set ID [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about the snapshot that has the specified data set ID.

``[-msid <integer>]`` - Snapshot Master Data Set ID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the snapshot that has the specified master data set ID.

``[-create-time <Date>]`` - Creation Time::
If this parameter is specified, the command displays information only
about the snapshots that match the specified creation time.

``[-busy {true|false}]`` - Snapshot Busy::
If this parameter is specified, the command displays information only
about the snapshots that have the specified busy status.
```

```
``[-owners <text>,...]`` - List of Owners::
If this parameter is specified, the command displays information only
about the snapshots that are owned by the specified list of owners.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Snapshot Size::
If this parameter is specified, the command displays information only
about the snapshots that have the specified size. The size is specified as
a character specifying the unit of measurement after a number specifying
the size in the mentioned unit of measurement: k (kilobytes), m
(megabytes), g (gigabytes), or t (terabytes). If the unit of measurement
is not specified, bytes are used as the unit, and the specified number is
rounded up to the nearest 4 KB. You may also use an inequality such as
``_>10 MB_`` as input.

``[-blocks <percent>]`` - Percentage of Total Blocks::
If this parameter is specified, the command displays information only
about the snapshots that have the specified percentage of total blocks on
their parent volumes. You may also use an inequality such as ``_>10_`` as
input.

``[-usedblocks <percent>]`` - Percentage of Used Blocks::
If this parameter is specified, the command displays information only
about the snapshots that have the specified percentage of used blocks on
their parent volumes. You may also use an inequality such as ``_>10_`` as
input.

``[-cpcount <integer>]`` - Consistency Point Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the snapshots that have the specified number of consistency points.
You may also use an inequality such as ``_<100_`` as input.

``[-comment <text>]`` - Comment::
If this parameter is specified, the command displays information only
about the snapshots that have the specified comment text. You may also
specify an inequality such as ``_!"-"_`` as input.

``[-fs-version <text>]`` - File System Version [.advaced-priv]#(privilege:
advanced)#::
If you use this parameter the only snapshots displayed are those that were
created when the file system was of a specific release. This parameter is
helpful especially when you need to upgrade to newer software release and
want to know the snapshots that will be impacted by the upgrade process.

``[-logical-snap-id <integer>]`` - Logical Snap ID [.advaced-
priv]#(privilege: advanced)#::
```

```
If you use this parameter only those snapshots which have the specified
logical snapshot ID will be shown. Logical snapshot ids are internal
implementation details of volume snapshots.

``[-is-7-mode {true|false}]`` - 7-Mode Snapshot::
If you use this parameter only those snapshots which have the specified
value are shown. This value is ``_true_`` for the snapshots that exist on
the volume that was in 7-mode configuration and then transitioned to a
clustered configuration. In such a scenario, the volume is in a clustered
configuration and the existing snapshots are still in the 7-mode
configuration.

``[-snapmirror-label <text>]`` - Label for SnapMirror Operations::
If you use this parameter, only those snapshots that have the specified
SnapMirror Label value are shown.

``[-state {valid|invalid|partial|pre-conversion}]`` - Snapshot State::
If you use this parameter only those snapshots which have the specified
state will be shown.

``[-is-constituent {true|false}]`` - Constituent Snapshot::
If you use this parameter, only those snapshots whose parent volume is a
constituent volume of a FlexGroup will be shown.

``[-node <nodename>]`` - Node [.advaced-priv]#(privilege: advanced)#::
If you use this parameter only those snapshots that are located on the
specified storage system are shown.

``[-inofile-version <integer>]`` - Snapshot Inofile Version [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about the snapshots whose inode files are at the specified version.

``[-expiry-time <MM/DD/YYYY HH:MM:SS>]`` - Expiry Time::
If you use this parameter only those snapshots that have the specified
expiry time are shown.

``[-compression-type {none|secondary|adaptive}]`` - Compression Type
[.advaced-priv]#(privilege: advanced)#::
If you use this parameter only those snapshots that have the specified
compression type are shown.

``[-snaplock-expiry-time {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm] | infinite}]``
- SnapLock Expiry Time::
If you use this parameter only those snapshots that have the specified
SnapLock expiry time are shown.
```

```
``[-application-io-size <Application I/O Size>]`` - Application IO Size
[.advaced-priv]#(privilege: advanced)#::
If you use this parameter only those snapshots that have the specified
application IO size are shown.

``[-is-qtree-caching-enabled {true|false}]`` - Is Qtree Caching Support
Enabled [.advaced-priv]#(privilege: advanced)#::
If you use this parameter, only those snapshots of FlexGroup volumes or
origin of a FlexCache volumes that have the specified Qtree caching status
are shown.

``[-snaplock-snapshot-expired {true|false}]`` - Has the SnapLock Snapshot
Expired::
If this parameter is specified, the command displays information as to
whether or not a locked snapshot has expired.

``[-seconds-until-snaplock-snapshot-expiry <integer>]`` - Seconds until
the SnapLock Snapshot Expiry::
If this parameter is specified, the command displays remaining SnapLock
expiry time (in seconds) of a locked snapshot.

== Examples

The following example displays default information about all snapshots of
a volume named vol1:
```

cluster1::> volume snapshot show -volume vol1
---Blocks---
Vserver Volume Snapshot Size Total% Used%
-------- -------- ---------- -------- ------ -----
cluster1 vol1
one 68KB 0% 33%
two 72KB 0% 34%
2 entries were displayed.

```
The following example displays snapshots which are older than 1 hour,
limiting the output to wanted fields:
```

cluster1::> volume snapshot show -create-time <1h -fields create-time, size
vserver volume snapshot create-time size
-------- ------ -------- ---------- ----
cluster1 vol1 one Mon Nov 17 10:23:42 2014 68KB
cluster1 vol1 two Mon Nov 17 10:23:44 2014 72KB
2 entries were displayed.

```
The following example displays detailed information about a specific
snapshot, using the 'snap' alias:
```

cluster1::> snap show -volume vol1 -snapshot one -instance
Vserver: cluster1
Volume: vol1
Snapshot: one
Snapshot Data Set ID: 4294968322
Snapshot Master Data Set ID: 6442451970
Creation Time: Mon Nov 17 10:23:42 2014
Snapshot Busy: false
List of Owners: -
Snapshot Size: 68KB
Percentage of Total Blocks: 0%
Percentage of Used Blocks: 33%
Consistency Point Count: 4
Comment: -
File System Version: 9.0
7-Mode Snapshot: false
Label for SnapMirror Operations: -
Constituent Snapshot: false
Node: node1
Snapshot Inofile Version: 3
Expiry Time: -
SnapLock Expiry Time: -
SnapLock Snapshot Expired: -
Seconds until SnapLock Snapshot Expiry: -

```
[[ID5b65503e0437f3c92b75d2ee25dc7c81]]
= volume snapshot autodelete modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify autodelete settings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot autodelete modify` command enables you to modify
snapshot autodelete and LUN, NVMe namespace or file clone autodelete
policy settings. Based on the defined policy, automatic deletion of
snapshots and LUN, NVMe namespace or file clones is triggered. Automatic
deletion of snapshots and LUN, NVMe namespace or file clones is useful
```

```
when you want to automatically reclaim space consumed by the snapshots and
LUN, NVMe namespace or file clones from the volume when it is low in
available space. LUN, NVMe namespace or file clone autodelete follows
snapshot autodelete. This command works only on a read-write parent
volume. You cannot setup automatic snapshot deletion and automatic LUN,
NVMe namespace or file clone deletion for read-only volumes.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the volume is located.

``-volume <volume name>`` - Volume Name::
This specifies the volume whose autodelete policy has to be modified.

``[-enabled {true|false}]`` - Enabled::
This option specifies whether automatic deletion of snapshots and LUN,
NVMe namespace or file clones is enabled or disabled. If set to ``_true_``
, automatic deletion of snapshots and LUN, NVMe namespace or file clones
is enabled. If set to ``_false_`` , automatic deletion of snapshots and
LUN, NVMe namespace or file clones is disabled.

``[-commitment {try|disrupt|destroy}]`` - Commitment::
This option specifies which snapshots and LUN, NVMe namespace or file
clones can be automatically deleted to reclaim back space. + When set to
``_try_`` , the snapshots which are not locked by any application and the
LUN, NVMe namespace or file clones which are not configured as preserved
are deleted. + When set to ``_disrupt_`` , the snapshots which are not
locked by data backing functionalities (such as volume clones, LUN clones,
NVMe namespace clones and file clones) and LUN, NVMe namespace or file
clones which are not configured as preserved are deleted. In the
``_disrupt_`` mode, the snapshots locked by data protection utilities such
as SnapMirror and Volume Move can be deleted. If such a locked snapshot is
deleted during the data transfer, the transfer is aborted. + When set to
``_destroy_`` , the snapshots locked by the data backing functionalities
are deleted. In addition, all the LUN, NVMe namespace or file clones in
the volume are deleted.

``[-defer-delete {scheduled|user_created|prefix|none}]`` - Defer Delete::
This option determines the order in which snapshots can be deleted. +
Possible values are as follows:
+
  * When set to ``_scheduled_`` , scheduled snapshots are the last to be
deleted.
  * When set to ``_user_created_`` , user snapshots are the last to be
deleted.
```

```
  * When set to ``_prefix_`` , snapshots matching a certain prefix are the
last to be deleted.
  * When set to ``_none_`` , no defer deletion order is honored.

+
 This option is not applicable for LUN, NVMe namespace or file clones.

``[-delete-order {newest_first|oldest_first}]`` - Delete Order::
This option specifies if the oldest snapshot and the oldest LUN, NVMe
namespace or file clone or the newest snapshot and the newest LUN, NVMe
namespace or file clone are deleted first.

``[-defer-delete-prefix <text>]`` - Defer Delete Prefix::
This option specifies the prefix string for the `-defer-delete prefix`
parameter. The option is not applicable for LUN, NVMe namespace or file
clones.

``[-target-free-space <percent>]`` - Target Free Space::
This option specifies the free space percentage at which the automatic
deletion of snapshots and LUN, NVMe namespace or file clones must stop.
Depending on the `-trigger` snapshots and LUN, NVMe namespace or file
clones are deleted until you reach the target free space percentage.

``[-trigger {volume|snap_reserve|(DEPRECATED)-space_reserve}]`` -
Trigger::
This option specifies the condition which starts the automatic deletion of
snapshots and LUN, NVMe namespace or file clones. + Setting this option to
``_volume_`` triggers automatic deletion of snapshots and LUN, NVMe
namespace or file clones when the volume reaches threshold capacity and
the volume space reserved for snapshots is exceeded. + Setting the option
to ``_snap_reserve_`` triggers automatic deletion of snapshots and LUN,
NVMe namespace or file clones when the space reserved for snapshots
reaches threshold capacity. + Setting the option to (DEPRECATED)-
``_space_reserve_`` triggers automatic deletion of snapshots when reserved
space in the volume reaches threshold capacity and the volume space
reserved for snapshots is exceeded. +
+
NOTE: The option ``_space_reserve_`` is deprecated.
+
The threshold capacity is determined by the size of the volume as follows:
+
  * If the volume size is less than 20 GB, the autodelete threshold is
85%.
  * If the volume size is equal to or greater than 20 GB and less than 100
GB, the autodelete threshold is 90%.
  * If the volume size is equal to or greater than 100 GB and less than
```

```
500 GB, the autodelete threshold is 92%.
  * If the volume size is equal to or greater than 500 GB and less than 1
TB, the autodelete threshold is 95%.
  * If the volume size is equal to or greater than 1 TB, the autodelete
threshold is 98%.



``[-destroy-list <text>]`` - Destroy List::
This option specifies a comma separated list of data backing functions
which are affected if the automatic deletion of the snapshot backing that
service is triggered. The possible values for this option are
``_lun_clone,fileclone_`` , ``_lun_clone,sfsr_`` , ``_vol_clone_`` ,
``_cifs_share_`` , or ``_none_`` . Except ``_none_`` , all the other
options can be combined as a comma separated list. Note that "lun_clone",
"file_clone" and "sfsr" individually are not valid values. Only pairs
"lun_clone,file_clone" and "lun_clone,sfsr" are supported.
+
NOTE: For the purposes of autodelete, ``_lun_clone_`` includes both LUNs
and NVMe namespaces.
+
If you specify ``_vol_clone_`` , the cloned volume backed by the snapshot
is deleted. + If you specify ``_lun_clone_`` , and a LUN or NVMe namespace
is in the process of being cloned when autodelete is triggered, the
cloning operation is aborted. Any access to this LUN or NVMe namespace
will result in an error being reported to the client. + If you specify
``_file_clone_`` , and the file cloning operation is in progress when
autodelete is triggered, the cloning operation is aborted. Any access to
this file will result in an error being reported to the client. + If you
specify ``_sfsr_`` , and the file restore is in progress when autodelete
is triggered, the restore operation is aborted. + If the snapshot is
locked either by a ``_lun_clone_`` or ``_file_clone_`` or both, the `-
destroy-list` must be set to ``_lun_clone,file_clone_`` . + If the
snapshot is locked either by a ``_lun_clone_`` or ``_sfsr_`` operation or
both, the `-destroy-list` must be set to ``_lun_clone,file_clone_`` . The
options ``_file_clone_`` and ``_sfsr_`` are equivalent to each other. + If
you set `-destroy-list` to ``_lun_clone,file_clone_`` and the snapshot is
backing a file clone or sfsr operation, both the operations are aborted.
This is also the case when you set `-destroy-list` to ``_lun_clone,sfsr_``
. + LUN, NVMe namespace or file clone autodelete is applicable only if `-
destroy-list` contains lun_clone and file_clone

== Examples

The following example enables snapshot autodelete and sets the trigger to
``_snap_reserve_`` for volume ``_vol3_`` which is part of the Vserver
```

```
``_vs0_`` :
```

cluster1::> volume snapshot autodelete modify -vserver vs0 -volume vol3 -enabled true -trigger snap_reserve

```
The following example enables snapshot autodelete and LUN, NVMe namespace
or file clone autodelete for volume ``_vol3_`` which is part of the
Vserver ``_vs0_`` :
```

cluster1::> volume snapshot autodelete modify -vserver vs0 -volume vol3 -enabled true -trigger volume -commitment try -delete-order oldest_first -destroy-list lun_clone,file_clone

```
[[ID925cb6b97d1ee6f4ffadaaf312a37b76]]
= volume snapshot autodelete show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display autodelete settings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot autodelete show` command displays information about
snapshot autodelete policies. The command output depends on the parameters
specified with the command. If no parameters are specified, the command
displays a table with the following information about all the available
snapshot autodelete policies:

* Vserver name
* Volume name
* Option name
* Option value

To display a detailed list view with additional information, run the
command and select the `-instance` view. The detailed list view provides
the following information:

* Vserver name
* Volume name
* Enabled
* Commitment
```

```
* Defer Delete
* Delete Order
* Defer Delete Prefix
* Target Free Space
* Trigger
* Destroy List
* Is Constituent Volume



By using the `-fields` option you can choose to print only the certain
fields in the output. This presents the selected fields in a table view.
This is ideal when you want additional information to be different from
the information that is provided by the default table view, but would like
it in a format which is visually easy to compare.

You can specify additional parameters to display the information that
matches only those parameters. For example, to display information only
about snapshot autodelete policies which are enabled, run the command with
`-enabled true` parameter. If you specify multiple filtering parameters,
only those policies that match all the specified parameters are displayed.

== Parameters

{ ``[-fields <fieldname>,...]``::
This option allows you to print only certain fields in the output.

| ``[-instance  ]`` }::
This option allows you to print a detailed list view.

``[-vserver <vserver name>]`` - Vserver Name::
If this parameter and the `-volume` parameter are specified, the command
displays detailed autodelete policy information about the specified
volume. If this parameter is specified by itself, the command displays
autodelete policy information about volumes on the specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If this parameter and the `-vserver` parameter are specified, the command
displays detailed autodelete policy information about the specified
volume. If this parameter is specified by itself, the command displays
autodelete policy information about all volumes matching the specified
name.

``[-enabled {true|false}]`` - Enabled::
If this parameter is specified, the command displays information about
autodelete policies that match the specified parameter value.
```

```
``[-commitment {try|disrupt|destroy}]`` - Commitment::
If this parameter is specified, the command displays information about
autodelete policies that match the specified commitment value.

``[-defer-delete {scheduled|user_created|prefix|none}]`` - Defer Delete::
If this parameter is specified, the command displays information about
autodelete policies that match the specified defer deletion criterion.

``[-delete-order {newest_first|oldest_first}]`` - Delete Order::
If this parameter is specified, the command displays information about
autodelete policies that match the specified deletion order.

``[-defer-delete-prefix <text>]`` - Defer Delete Prefix::
If this parameter is specified, the command displays information about
autodelete policies that match the prefix used for deferring deletion.

``[-target-free-space <percent>]`` - Target Free Space::
If this parameter is specified, the command displays information about
autodelete policies that match the specified target free space.

``[-trigger {volume|snap_reserve|(DEPRECATED)-space_reserve}]`` -
Trigger::
If this parameter is specified, the command displays information about
autodelete policies that match the specified trigger condition.

``[-destroy-list <text>]`` - Destroy List::
If this parameter is specified, the command displays information about
autodelete policies that match the specified value.

``[-is-constituent {true|false}]`` - Is Constituent Volume::
If this parameter is specified, the command displays information about
autodelete policies for the constituent volumes of FlexGroup volumes.

== Examples

The following example displays snapshot autodelete policy settings for
volume vol3 which is inside the Vserver vs0:
```

cluster1::> volume snapshot autodelete show -vserver vs0 -volume vol3

Vserver Volume Option Name Option Value
--------- ---------- ---------- ----------
vs0 vol3
Enabled false
Commitment try
Trigger volume
Target Free Space 20%

Delete Order oldest_first
Defer Delete user_created
Defer Delete Prefix (not specified)
Destroy List none

```
[[IDf5955285d61290b77278e7d2e204448f]]
= volume snapshot policy add-schedule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a schedule to snapshot policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot policy add-schedule` command adds a schedule to a
snapshot policy. You can create a schedule by using the
xref:{relative_path}job-schedule-cron-create.html[job schedule cron
create] or xref:{relative_path}job-schedule-interval-create.html[job
schedule interval create] commands.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which a snapshot policy schedule is to be
added.

``-policy <snapshot policy>`` - Snapshot Policy Name::
This specifies the snapshot policy to which a schedule is to be added.

``-schedule <text>`` - Schedule Name::
This specifies the schedule that is to be added to the snapshot policy.

``-count <integer>`` - Maximum Snapshot Copies for Schedule::
This specifies the maximum number of snapshots that can be taken by the
specified schedule. The total count of all the snapshots to be retained
for the policy cannot be more than 1023.

``[-prefix <text>]`` - Snapshot Copy Name Prefix for Schedule::
This option specifies the prefix with which snapshots will be created for
the added schedule. Every schedule has only one prefix. Once a prefix gets
associated with a schedule, you cannot update the prefix. If some prefix
```

```
is already associated with the schedule and you do not specify this
parameter, then the previously defined prefix is used. The command fails
if you try to update an existing prefix for a schedule. If no prefix is
associated with the schedule and you do not specify this parameter, then
schedule name is be used as the prefix.

``[-snapmirror-label <text>]`` - Label for SnapMirror Operations::
This specifies the SnapMirror Label identified with a snapshot when it is
created for the added schedule. The SnapMirror Label is used by the
Vaulting subsystem when you back up snapshots to the Vault Destination.

``[-retention-period <snaplock minmax period>]`` - SnapLock Retention
Period::
This specifies the retention period for snapshots created by the schedule.

== Examples

The following example adds a schedule named midnight to the snapshot
policy named snappolicy_nightly on Vserver vs0. The schedule can take a
maximum of five snapshots.
```

cluster1::> volume snapshot policy add-schedule -vserver vs0 -policy snappolicy_nightly -schedule midnight -count 5 -retention-period "7 days"

```
== Related Links

* xref:{relative_path}job-schedule-cron-create.html[job schedule cron
create]
* xref:{relative_path}job-schedule-interval-create.html[job schedule
interval create]




[[ID25313fa8c33fcca3027df28b407fcc21]]
= volume snapshot policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new snapshot policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `volume snapshot policy create` command creates a snapshot policy. A
snapshot policy includes at least one schedule, up to a maximum of five
schedules, and a maximum number of snapshots per schedule. You can create
a schedule by using the xref:{relative_path}job-schedule-cron-
create.html[job schedule cron create] or xref:{relative_path}job-schedule-
interval-create.html[job schedule interval create] commands. When applied
to a volume, the snapshot policy specifies the schedule on which snapshots
are taken and the maximum number of snapshots that each schedule can take.
The total count of all the snapshots to be retained for the policy cannot
be more than 1023.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the snapshot policy is to be created.

``-policy <snapshot policy>`` - Snapshot Policy Name::
This specifies the snapshot policy that is to be created.

``-enabled {true|false}`` - Snapshot Policy Enabled::
This specifies whether the snapshot policy is enabled.

``[-comment <text>]`` - Comment::
This option specifies a text comment for the snapshot policy.

``-schedule1 <text>`` - Schedule1 Name::
This specifies the name of the first schedule associated with the snapshot
policy.

``-count1 <integer>`` - Maximum Snapshot Copies for Schedule1::
This specifies the maximum number of snapshots that can be taken by the
first schedule.

``[-prefix1 <text>]`` - Snapshot Copy Name Prefix for Schedule1::
This option specifies the prefix associated with the first schedule. Every
schedule has only one prefix. The command fails if you try to update an
existing prefix. If you do not specify this parameter and there is no
prefix associated with the schedule, the schedule name is used as the
prefix. If you do not specify this parameter and there is already a prefix
associated with the schedule from a previous invocation of the command,
then that prefix is used.

``[-snapmirror-label1 <text>]`` - Label for SnapMirror Operations for
Schedule1::
```

```
This specifies the SnapMirror Label of the first schedule associated with
the snapshot policy. Once specified, all snapshots created for that
schedule have the SnapMirror Label assigned to them. The SnapMirror Label
is used by the Vaulting subsystem when you back up snapshots to the Vault
Destination.

``[-retention-period1 <snaplock minmax period>]`` - SnapLock Retention
Period::
This specifies the retention period for snapshots that can be taken by the
first schedule. If this parameter is specified, the `-schedule1` parameter
must also be specified. The default is no retention period.

``[-schedule2 <text>]`` - Schedule2 Name::
This option specifies the name of the second schedule associated with the
snapshot policy. If this parameter is specified, the `-count2` parameter
must also be specified.

``[-count2 <integer>]`` - Maximum Snapshot Copies for Schedule2::
This option specifies the maximum number of snapshots that can be taken by
the second schedule. If this parameter is specified, the `-schedule2`
parameter must also be specified.

``[-prefix2 <text>]`` - Snapshot Copy Name Prefix for Schedule2::
This option specifies the prefix associated with the second schedule. If
this parameter is specified, `-schedule2` and `-count2` parameters must
also be specified. Every schedule has only one prefix. The command fails
if you try to update an existing prefix. If you do not specify this
parameter and there is no prefix associated with the schedule, the
schedule name is used as the prefix. If you do not specify this parameter
and there is already a prefix associated with the schedule from a previous
invocation of the command, then that prefix is used.

``[-snapmirror-label2 <text>]`` - Label for SnapMirror Operations for
Schedule2::
This specifies the SnapMirror Label of the second schedule associated with
the snapshot policy. Once specified, all snapshots created for that
schedule have the SnapMirror Label assigned to them. The SnapMirror Label
is used by the Vaulting subsystem when you back up snapshots to the Vault
Destination.

``[-retention-period2 <snaplock minmax period>]`` - SnapLock Retention
Period::
This specifies the retention period for snapshots that can be taken by the
second schedule. If this parameter is specified, the `-schedule2`
parameter must also be specified. The default is no retention period.
```

```
``[-schedule3 <text>]`` - Schedule3 Name::
This option specifies the name of the third schedule associated with the
snapshot policy. If this parameter is specified, the `-count3` parameter
must also be specified.

``[-count3 <integer>]`` - Maximum Snapshot Copies for Schedule3::
This option specifies the maximum number of snapshots that can be taken by
the third schedule. If this parameter is specified, the `-schedule3`
parameter must also be specified.

``[-prefix3 <text>]`` - Snapshot Copy Name Prefix for Schedule3::
This option specifies the prefix associated with the third schedule. If
this parameter is specified, `-schedule3` and `-count3` parameters must
also be specified. Every schedule has only one prefix. The command fails
if you try to update an existing prefix. If you do not specify this
parameter and there is no prefix associated with the schedule, the
schedule name is used as the prefix. If you do not specify this parameter
and there is already a prefix associated with the schedule from a previous
invocation of the command, then that prefix is used.

``[-retention-period3 <snaplock minmax period>]`` - SnapLock Retention
Period::
This specifies the retention period for snapshots that can be taken by the
third schedule. If this parameter is specified, the `-schedule3` parameter
must also be specified. The default is no retention period.

``[-snapmirror-label3 <text>]`` - Label for SnapMirror Operations for
Schedule3::
This specifies the SnapMirror Label of the third schedule associated with
the snapshot policy. Once specified, all snapshots created for that
schedule have the SnapMirror Label assigned to them. The SnapMirror Label
is used by the Vaulting subsystem when you back up snapshots to the Vault
Destination.

``[-schedule4 <text>]`` - Schedule4 Name::
This option specifies the name of the fourth schedule associated with the
snapshot policy. If this parameter is specified, the `-count4` parameter
must also be specified.

``[-count4 <integer>]`` - Maximum Snapshot Copies for Schedule4::
This option specifies the maximum number of snapshots that can be taken by
the fourth schedule. If this parameter is specified, the `-schedule4`
parameter must also be specified.

``[-prefix4 <text>]`` - Snapshot Copy Name Prefix for Schedule4::
This option specifies the prefix associated with the fourth schedule. If
```

this parameter is specified, `-schedule4` and `-count4` parameters must also be specified. Every schedule has only one prefix. The command fails if you try to update an existing prefix. If you do not specify this parameter and there is no prefix associated with the schedule, the schedule name is used as the prefix. If you do not specify this parameter and there is already a prefix associated with the schedule from a previous invocation of the command, then that prefix is used.

``[-retention-period4 <snaplock minmax period>]`` - SnapLock Retention Period::
This specifies the retention period for snapshots that can be taken by the fourth schedule. If this parameter is specified, the `-schedule4` parameter must also be specified. The default is no retention period.

``[-snapmirror-label4 <text>]`` - Label for SnapMirror Operations for Schedule4::
This specifies the SnapMirror Label of the fourth schedule associated with the snapshot policy. Once specified, all snapshots created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up snapshots to the Vault Destination.

``[-schedule5 <text>]`` - Schedule5 Name::
This option specifies the name of the fifth schedule associated with the snapshot policy. If this parameter is specified, the `-count5` parameter must also be specified.

``[-count5 <integer>]`` - Maximum Snapshot Copies for Schedule5::
This option specifies the maximum number of snapshots that can be taken by the fifth schedule. If this parameter is specified, the `-schedule5` parameter must also be specified.

``[-prefix5 <text>]`` - Snapshot Copy Name Prefix for Schedule5::
This option specifies the prefix associated with the fifth schedule. If this parameter is specified, `-schedule5` and `-count5` parameters must also be specified. Every schedule has only one prefix. The command fails if you try to update an existing prefix. If you do not specify this parameter and there is no prefix associated with the schedule, the schedule name is be used as the prefix. If you do not specify this parameter and there is already a prefix associated with the schedule from a previous invocation of the command, then that prefix is used.

``[-retention-period5 <snaplock minmax period>]`` - SnapLock Retention Period::
This specifies the retention period for snapshots that can be taken by the fifth schedule. If this parameter is specified, the `-schedule5` parameter

```
must also be specified. The default is no retention period.

``[-snapmirror-label5 <text>]`` - Label for SnapMirror Operations for
Schedule5::
This specifies the SnapMirror Label of the fifth schedule associated with
the snapshot policy. Once specified, all snapshots created for that
schedule have the SnapMirror Label assigned to them. The SnapMirror Label
is used by the Vaulting subsystem when you back up snapshots to the Vault
Destination.

== Examples

The following example creates a snapshot policy named snappolicy_4hrs on a
Vserver named vs0. The policy runs on a single schedule named 4hrs with a
prefix every_4_hour and has a maximum number of five snapshots.
```

cluster1::> volume snapshot policy create -vserver vs0 -policy snappolicy_4hrs -schedule1 4hrs -count1 5 -prefix1 every_4_hour -retention-period1 "8 hours"

```
== Related Links

* xref:{relative_path}job-schedule-cron-create.html[job schedule cron
create]
* xref:{relative_path}job-schedule-interval-create.html[job schedule
interval create]




[[IDc481152c7cad02b55d03a13f446af89b]]
= volume snapshot policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a snapshot policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot policy delete` command deletes a snapshot policy.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the snapshot policy is to be deleted.

``-policy <snapshot policy>`` - Snapshot Policy Name::
This specifies the snapshot policy that is to be deleted.

== Examples

The following example deletes a snapshot policy named snappolicy_hourly on
Vserver vs0:
```

cluster1::> volume snapshot policy delete -vserver vs0 -policy snappolicy_hourly

```
[[ID4da40a325f86557d7043d711ac0ab593]]
= volume snapshot policy modify-schedule
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a schedule within snapshot policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot policy modify-schedule` command modifies the maximum
number of snapshots that can be taken by a snapshot policy's schedule.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which a snapshot policy schedule is to be
modified.

``-policy <snapshot policy>`` - Snapshot Policy Name::
This specifies the snapshot policy whose schedule is to be modified.

``-schedule <text>`` - Schedule Name::
This specifies the schedule that is to be modified.

``[-newcount <integer>]`` - Maximum Snapshot Copies for Schedule::
This specifies the maximum number of snapshots that can be taken by the
specified schedule. The total count of all the snapshots to be retained
for the policy cannot be more than 1023.

``[-newsnapmirror-label <text>]`` - Label for SnapMirror Operations::
This specifies the SnapMirror Label identified with a snapshot when it is
created for the specified schedule. The SnapMirror Label is used by the
Vaulting subsystem when you back up snapshots to the Vault Destination. If
an empty label ("") is specified, the existing label will be deleted.

``[-newretention-period <snaplock minmax period>]`` - SnapLock Retention
Period::
This specifies the retention period for snapshots created by the schedule.

== Examples

The following example changes the maximum number of snapshots from five to
four for a schedule named midnight on a snapshot policy named
snappolicy_nightly on Vserver vs0:
```

cluster1::> volume snapshot policy modify-schedule -vserver vs0 -policy snappolicy_nightly -schedule midnight -newcount 4 -newretention-period "7 days"

```
[[ID566027dc5a4336472879ed0691496d7b]]
= volume snapshot policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a snapshot policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot policy modify` command enables you to modify the
description associated with a snapshot policy and whether the policy is
enabled or disabled.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which the snapshot policy is to be modified.

``-policy <snapshot policy>`` - Snapshot Policy Name::
This specifies the snapshot policy that is to be modified.

``[-enabled {true|false}]`` - Snapshot Policy Enabled::
This optionally specifies whether the snapshot policy is enabled.

``[-comment <text>]`` - Comment::
This specifies the comment text for the snapshot policy.

``[-snapmirror-labels <text>,...]`` - Label for SnapMirror Operations::
This optionally specifies a comma separated list of SnapMirror labels that
are applied to the schedules in the snapshot policy. Each label in the
list applies to only one schedule in the snapshot policy (maximum of 5
SnapMirror labels), the first label applying to the first schedule, the
second label applying to the second schedule, and so on. You can have a
maximum of five SnapMirror labels, which corresponds to the maximum number
of schedules a snapshot policy can have. If an empty string ("") is
specified, the existing labels will be deleted from all the schedules.
```

```
== Examples

The following example changes the description of a snapshot policy named
snappolicy_wknd on Vserver vs0 to "Runs only on weekends":
```

cluster1::> volume snapshot policy modify -vserver vs0 -policy snappolicy_wknd -comment "Runs only on weekends"

```
[[ID26e31ece14037668a78cb176d791f003]]
= volume snapshot policy remove-schedule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a schedule from snapshot policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot policy remove-schedule` command removes a schedule
from a snapshot policy.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This specifies the Vserver on which a snapshot policy schedule is to be
removed.

``-policy <snapshot policy>`` - Snapshot Policy Name::
This specifies the snapshot policy from which a schedule is to be removed.

``-schedule <text>`` - Schedule Name::
This specifies the schedule that is to be removed from the snapshot
policy.

== Examples

The following example removes a schedule named hourly from a snapshot
policy named snappolicy_daily on Vserver vs0:
```

cluster1::> volume snapshot policy remove-schedule -vserver vs0 -policy snappolicy_daily -schedule hourly

```
[[ID534620199584270376584299dee04641]]
= volume snapshot policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show snapshot policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `volume snapshot policy show` command displays the following
information about snapshot policies:

* Vserver name
* Snapshot policy name
* Number of schedules in the policy
* Comment for the policy
* Individual schedule names
* Maximum number of snapshots associated with each schedule
* Snapshot name prefixes for the schedules
* SnapMirror Labels associated with the schedules



== 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.

| ``[-revert-incompatible  ]`` [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays snapshot policies
that are not supported in ONTAP 8.2. The total snapshot count in the
policy needs to be reduced to be equal to or less than the supported count
for the revert operation to succeed.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name::
```

```
If this parameter is specified, the command displays snapshot policies on
the specified Vserver.

``[-policy <snapshot policy>]`` - Snapshot Policy Name::
If this parameter is specified, the command displays detailed information
about the specified snapshot policy.

``[-enabled {true|false}]`` - Snapshot Policy Enabled::
If this parameter is specified, the command displays detailed information
only about the snapshot policy or policies that have the specified enabled
value.

``[-comment <text>]`` - Comment::
If this parameter is specified, the command displays information only
about the snapshot policy or policies that have the specified comment.

``[-total-schedules <integer>]`` - Total Number of Schedules in This
Policy::
If this parameter is specified, the command displays information only
about the snapshot policy or policies that have the specified total number
of schedules.

``[-schedules <text>,...]`` - Schedule Name::
If this parameter is specified, the command displays information only
about the snapshot policy or policies that have the specified list of
schedules.

``[-counts <integer>,...]`` - Maximum Snapshots for the Schedule::
If this parameter is specified, the command displays information only
about the snapshot policy or policies that have the specified list of
maximum numbers of snapshots per schedule.

``[-prefixes <text>,...]`` - Prefix Name::
If this parameter is specified, the command displays information only
about the snapshot policy or policies that have the specified list of
prefixes.

``[-snapmirror-labels <text>,...]`` - Label for SnapMirror Operations::
If this parameter is specified, the command displays information only
about the snapshot policies that have the specified SnapMirror Label. When
you specify a list of SnapMirror labels, the command displays all the
snapshot policies that contain any of the SnapMirror Labels specified in
the list.

``[-policy-owner <text>]`` - Owner of the policy::
If this parameter is specified, the command displays information only
about the snapshot policies that have the specified policy owner.
```

```
``[-total-count <integer>]`` - Total Number of Snapshots in This Policy::
If this parameter is specified, the command displays information only
about the snapshot policies that have the specified total number of
snapshots.

``[-retention-periods <snaplock minmax period>,...]`` - SnapLock Retention
Period::
If this parameter is specified, the command displays information only
about the snapshot policy or policies that have the specified retention-
period set.

== Examples

The following example displays information about all snapshot policies:
```

cluster1::> volume snapshot policy show
Vserver: cm
Number of Is
Policy Name Schedules Enabled Comment
---------- --------- ------- ----------
default 3 false Default policy with hourly, daily weekly schedules.
Schedule Count Prefix SnapMirror Label
---------- ----- ---------- ----------
hourly 6 hourly -
daily 2 daily -
weekly 2 weekly -

default-1weekly 3 false Default policy with 6 hourly, 2 daily 1 weekly schedule.
Schedule Count Prefix SnapMirror Label
---------- ----- ---------- ----------
hourly 6 hourly -
daily 2 daily -
weekly 1 weekly -

none 0 false Policy for no automatic snapshots.
Schedule Count Prefix SnapMirror Label
---------- ----- ---------- ----------
- - - -

Vserver: vs0
Number of Is
Policy Name Schedules Enabled Comment
---------- --------- ------- ----------
p1 1 false -
Schedule Count Prefix SnapMirror Label
---------- ----- ---------- ----------
weekly 2 weekly -

p2 2 true -
Schedule Count Prefix SnapMirror Label

--- --- --- ---
hourly 6 hourly -
daily 2 daily -

5 entries were displayed.

```
:leveloffset: -1


= volume upgrade-tasks commands

:leveloffset: +1


[[ID9ed8e6cfc9b43f317f78e796685a6537]]
= volume upgrade-tasks run
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Triggers the M-Host upgrade tasks for the volume

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `volume upgrade-tasks run` command triggers the M-host upgrade tasks
for the volume and sets the volume upgrade status.

== Parameters

``-vserver <vserver name>`` - Vserver where the volume is located
[.advaced-priv]#(privilege: advanced)#::
This specifies the vserver on which the volume is located.

``-volume <volume name>`` - Volume on which we are running triggering the
upgrade [.advaced-priv]#(privilege: advanced)#::
This specifies the volume on which we have to trigger the upgrade-tasks.

== Examples

The following example triggers upgrade-tasks on volume named vol1 on a
vserver named vs0:
```

```
node::> volume upgrade-tasks run -vserver vs1 -volume vol1
```

```
:leveloffset: -1


:leveloffset: -1


= vserver commands

:leveloffset: +1


[[ID8bf0822a77809f5b4b9cdedb47cdb759]]
= vserver add-aggregates
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add aggregates to the Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver add-aggregates` command adds aggregates to the Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
Specifies the Vserver for which aggregates have to be added.

``-aggregates <aggregate name>,...`` - List of Aggregates to Be Added::
Specifies the list of aggregates to add to the Vserver. The root
aggregates should not be specified in this list because though the command
will return success, volumes cannot be created on root aggregates. In a
MetroCluster configuration, this command does not honor the remote
cluster's aggregates.

== Examples

The following example illustrates how to add aggregates ``_aggr1_`` and
``_aggr2_`` to a Vserver named ``_vs.example.com_`` :
```

cluster1::> vserver add-aggregates -vserver vs.example.com -aggregates aggr1,aggr2

```
[[IDcfe798e2b041fff770aae26e6bce7a40]]
= vserver add-protocols
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add protocols to the Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver add-protocols` command adds given protocols to a specified
Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
This specifies the Vserver that is to be modified.

``-protocols {nfs|cifs|fcp|iscsi|ndmp|nvme|s3}`` - Protocols::
This parameter specifies the list of protocols to be allowed to run on the
Vserver. Possible values include nfs, cifs, fcp, iscsi, ndmp and nvme.

== Examples

The following example shows adding protocol 'cifs' to a vserver named
vs0.example.com.
```

cluster1::> vserver add-protocols -vserver vs0.example.com -protocols cifs

```
[[IDefb0737397c2583d84e5ae09ca7f6ead]]
= vserver context
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set Vserver context

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

Cluster administrators can use the `vserver context` command to login to a
specified Vserver with a specified Vserver user name. All subsequent
commands will be issued in the context of that Vserver. The role of the
cluster administrator will be the same as that of the user name with which
the Vserver context was set. The context is valid for the duration of the
CLI or Web UI session in which it is specified. The
xref:{relative_path}exit.html[exit] command can be used to return to the
original context.

== Parameters

``-vserver <vserver>`` - Vserver::
Use this parameter to specify the Vserver.

``[-username <text>]`` - Vserver Administrator User Name::
Use this parameter to specify a Vserver administrator user name for the
context. The default value ``_vsadmin_`` is used if one is not specified.

== Examples

The following example sets the CLI context to Vserver
``_vs0.example.com_`` . All subsequently issued commands will be executed
in the context of that Vserver:
```

cluster1::> vserver context -vserver vs0.example.com
Info: Use 'exit' command to return.
vs0.example.com::>

```
== Related Links
```

```
* xref:{relative_path}exit.html[exit]




[[ID79acb1eb702bf78ba183cfc8b26bc635]]
= vserver create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver create` command creates a Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
This specifies the name of the Vserver that is to be created. Use a fully
qualified domain name (FQDN) - for example, "data.example.com" - for the
Vserver to ensure unique Vserver names across cluster leagues.
+
NOTE: Maximum number of characters supported is 47, and 41 for a Vserver
with subtype "sync-source". "all" is a reserved name and must not be used
as a Vserver name.
+


``[-subtype <vserver subtype>]`` - Vserver Subtype::
This specifies the subtype of the Vserver being created. Possible values
are:
+
  * default - For default data Vservers
  * dp-destination - For Data Protection destination Vservers
  * sync-source - For MetroCluster source Vservers
  * sync-destination - For MetroCluster destination Vservers



``[-rootvolume <volume name>]`` - Root Volume::
```

```
This parameter optionally specifies the name of the Vserver's root volume,
which is created when the Vserver is created. The default name is
``_vserverName_root_`` . The size of the Vserver's root volume is 1GB

``[-aggregate <aggregate name>]`` - Aggregate::
This parameter optionally specifies the storage aggregate that holds the
Vserver's root volume. Selection of the aggregate is based on the Vserver
setup algorithm.
+
  * Creating a root volume on the SnapLock aggregate is not supported.
  * Creating a root volume of sync-source Vserver on the unmirrored
aggregate is not supported.



``[-rootvolume-security-style <security style>]`` - Root Volume Security
Style::
This parameter optionally specifies the security style for the Vserver's
root volume. Possible values include ``_unix_`` (for UNIX mode bits),
``_ntfs_`` (for CIFS ACLs), and ``_mixed_`` (for mixed NFS and CIFS
access). The default value is ``_unix_`` . Regardless of the security
style, both NFS and CIFS clients can read from and write to the root
volume.

``[-language <Language code>]`` - Default Volume Language Code::
This optionally specifies the default language encoding setting for the
Vserver and its volumes. The recommended format is to append ``_.UTF-8_``
for the language encoding values. For example, for the ``_en_US_``
language, the recommended format is ``_en_US.UTF-8_`` . The default
setting is ``_C.UTF-8_`` .

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
This optionally specifies the snapshot policy for new volumes created on
the Vserver. If no value is specified, the default snapshot policy is
used. You can use the `-snapshot-policy` parameter on the
xref:{relative_path}volume-create.html[volume create] or
xref:{relative_path}volume-modify.html[volume modify] commands to set the
snapshot policy on a specific volume, regardless of its Vserver's snapshot
policy setting.

``[-data-services <LIF Service Name>,...]`` - Data Services::
This optionally specifies the data services for new network interfaces
created on the Vserver. If no value is specified, the default services
list will be applied. This information will be used to construct the
default service policies for this Vserver, which can be viewed using the
`network interface service-policy show` command.
```

```
``[-comment <text>]`` - Comment::
This optionally specifies a comment for the Vserver.

``[-quota-policy <text>]`` - Quota Policy::
This optionally specifies a quota policy for the Vserver.

``[-caching-policy <text>]`` - Caching Policy Name::
 This optionally specifies the caching policy to apply to the Vserver. A
caching policy defines how the system caches this volume's data in Flash
Cache modules. If a caching policy is not assigned to this Vserver, the
system uses the default cluster-wide policy. The available caching
policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read, and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read, and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.

+
Default caching-policy is auto.

``[-ipspace <IPspace>]`` - IPspace Name::
This optionally specifies the IPspace the Vserver will be assigned to. If
left unspecified, the Vserver will be assigned to the default IPspace.

``[-foreground {true|false}]`` - Foreground Process::
This parameter optionally specifies whether the Vserver create operation
can be executed in the background. If nothing is specified, by default the
Vserver create operation is executed in the foreground.

``[-is-space-reporting-logical {true|false}]`` - Logical Space Reporting::
This optionally specifies whether to report space logically on residing
volumes. When space is reported logically, ONTAP reports the volume space
such that all the physical space saved by the storage efficiency features
are also reported as used.
```

```
``[-is-space-enforcement-logical {true|false}]`` - Logical Space
Enforcement::
This optionally specifies whether to perform logical space accounting on
residing volumes. When space is enforced logically, ONTAP enforces volume
settings such that all the physical space saved by the storage efficiency
features will be calculated as used.

``[-anti-ransomware-default-volume-state {disabled|dry-run}]`` - Default
Anti_ransomware State of the Vserver's Volumes::
This optionally specifies the default anti_ransomware state for volumes in
the Vserver. The available anti_ransomware states for the volumes in the
Vserver are:
+
  * disabled - Anti_ransomware protection is not enabled on the volumes in
the Vserver.
  * dry-run - Anti_ransomware protection is in learning mode on the
volumes in the Vserver.

+
Default anti-ransomware-default-volume-state is ``_disabled_`` .

``[-auto-enable-analytics {true|false}]`` - Enable Analytics on New
Volumes::
This parameter specifies whether analytics is automatically enabled on
volumes that will be created in the Vserver. The default value is false.

``[-auto-enable-activity-tracking {true|false}]`` - Enable Activity
Tracking on New Volumes::
This parameter specifies whether activity tracking is automatically
enabled on volumes that will be created in the Vserver. The default value
is false.

``[-storage-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Storage Limit::
This parameter specifies the maximum size limit for a Vserver. Default
value is 0, which means no limit.

``[-storage-limit-threshold-alert <percent>]`` - Storage Limit Threshold
Alert::
This parameter specifies when the "vserver.storage.threshold" event log
message is generated. This is a percentage value of storage-limit. Default
value is 90%.

``[-anti-ransomware-auto-switch-from-learning-to-enabled {true|false}]`` -
Anti-ransomware Auto-switch from Learning to Enabled::
This parameter optionally specifies whether anti-ransomware state of the
volumes in this Vserver are automatically switched by the system from
```

“learning” (dry-run) to “enabled” (active) state after sufficient learning. The default value is ``_true_`` .

``[-anti-ransomware-auto-switch-minimum-incoming-data-percent <percent>]`` - Anti-ransomware Auto-switch Minimum Incoming Data (in percentage)::
One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this Vserver from “learning” (dry-run) to “enabled” is that the volume should have sufficient data ingested to do the learning. This parameter optionally specifies the minimum amount of data (in percentage) to be written to a given volume during the learning period to automatically switch the anti-ransomware state from “learning” to “enabled”. The amount of data considered as ingested also includes the data that is deleted or overwritten after ingestion. The default value is 5%.

``[-anti-ransomware-auto-switch-duration-without-new-file-extension <integer>]`` - Anti-ransomware Auto-switch Duration Without New File Extension (in Days)::
One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this Vserver from “learning” (dry-run) to “enabled” is that no new file-extensions are observed in the volume in recent time. This parameter optionally specifies the recent time duration (in days) to be considered during which no new file-extension should be observed in a given volume to automatically switch the anti-ransomware state from ‘learning” to “enabled”. The default value is ``_3_`` days.

``[-anti-ransomware-auto-switch-minimum-learning-period <integer>]`` - Anti-ransomware Auto-switch Minimum Learning Period::
One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this Vserver from “learning” (dry-run) to “enabled” is that the volume should be in “learning” state for sufficient time period. This parameter optionally specifies the minimum number of days a given volume should be in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. The default value is ``_10_`` days.

``[-anti-ransomware-auto-switch-minimum-file-count <integer>]`` - Anti-ransomware Auto-switch Minimum File Count::
One of the many conditions to be satisfied to automatically switch the anti-ransomware state of the volumes in this Vserver from “learning” (dry-run) to “enabled” is that the volume should have sufficient number of files created in “learning” state. This parameter optionally specifies the minimum number of new files to be created in a given volume in “learning” state to automatically switch the anti-ransomware state from “learning” to “enabled”. The default value is ``_200_`` .

```
``[-anti-ransomware-auto-switch-minimum-file-extension <integer>]`` -
Anti-ransomware Auto-switch Minimum File Extension::
One of the many conditions to be satisfied to automatically switch the
anti-ransomware state of the volumes in this Vserver from “learning” (dry-
run) to “enabled” is that the volume should have sufficient unique file
extension count in “learning” state. This parameter optionally specifies
the minimum number of unique file extension count in a given volume in
“learning” state to automatically switch the anti-ransomware state from
“learning” to “enabled”. The default value is ``_10_`` days.

== Examples

The following example creates a Vserver named ``_vs0.example.com_`` in the
IPspace ipspace123. The Vserver's root volume is named ``_root_vs0_`` and
is located on aggregate ``_aggr0_`` . The Vserver uses NIS for network
information, a file for name mapping information, and the language is U.S.
English:
```

cluster1::> vserver create -vserver vs0.example.com -ipspace ipspace123 -rootvolume root_vs0 -aggregate aggr0
-language en_US.UTF-8 -rootvolume-security-style mixed

```
The following example creates a Vserver named ``_vs1_`` using default
values. The default name for the Vserver's root volume is
``_vserverName_root_`` and the Vserver is located on an aggregate selected
on the basis of the Vserver setup algorithm. The default root volume's
security style is set to ``_unix_`` .
```

cluster1::> vserver create -vserver vs1
cluster1::> vserver show -vserver vs1 -fields rootvolume, rootvolume-security-style, aggregate
vserver rootvolume aggregate rootvolume-security-style
------- ---------- --------- ------------------------
vs1 vs1_root aggr1 unix

```
== Related Links

* xref:{relative_path}volume-create.html[volume create]
* xref:{relative_path}volume-modify.html[volume modify]




[[ID683ec2c84118f02c9871c30594aeb867]]
= vserver delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver delete` command deletes a specified Vserver. If the Vserver
is associated with one or more volumes, you must manually delete volumes
(including root and mirror volumes) before you delete the Vserver. If the
Vserver subtype is ``_dp-destination_`` , change the Vserver subtype to
``_default_`` by specifying the Vserver as the destination in the
xref:{relative_path}snapmirror-break.html[snapmirror break] command before
deleting the objects owned by the Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
This specifies the Vserver that is to be deleted.

``[-foreground {true|false}]`` - Foreground Process::
This optionally specifies the Vserver delete operation can be executed in
the background. If nothing is specified, by default the Vserver delete
operation is executed in the foreground.

== Examples

The following example deletes a Vserver named ``_vs2.example.com_`` :
```

cluster1::> vserver delete -vserver vs2.example.com

```
== Related Links

* xref:{relative_path}snapmirror-break.html[snapmirror break]




[[ID5105b74ef640f45ee9f9a6f3094f07af]]
= vserver modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver modify` command modifies the attributes of a specified
Vserver. If the Vserver subtype is of type ``_dp-destination_`` , then
only the `-aggr-list` parameter can be modified.

== Parameters

``-vserver <vserver>`` - Vserver::
This specifies the Vserver that is to be modified.

``[-language <Language code>]`` - Default Volume Language Code::
This optional parameter specifies the default language encoding setting
for the Vserver and its volumes. The recommended format is to append
``_.UTF-8_`` for the language encoding values. For example, for the en_US
language, the recommended format is ``_en_US.UTF-8_`` . The default
setting is ``_C.UTF-8_`` .

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
This optional parameter specifies the snapshot policy for new volumes
created on the Vserver. Changing the snapshot policy for a Vserver will
not change the snapshot policy for the existing volumes within the Vserver
and it will only be applicable for the new volumes that will be created
later. You can use the -snapshot-policy parameter with the volume create
```

```
or volume modify commands to set the snapshot policy on a specific volume,
regardless of its Vserver's snapshot policy setting.

``[-comment <text>]`` - Comment::
This optional parameter specifies a comment for the Vserver.

``[-quota-policy <text>]`` - Quota Policy::
This optional parameter specifies a quota policy to be used for all
volumes associated with a Vserver. You can create and configure multiple,
different quota policies, but each Vserver must have one and only one
associated quota policy.

``[-aggr-list <aggregate name>,...]`` - List of Aggregates Assigned::
This optional parameter specifies a confined list of aggregates on which
volumes can be created for a Vserver by the Vserver administrator. But
these aggregates do not become exclusive property of the Vserver, i.e.
they might be assigned for use to other Vservers. If the value of this
parameter is specified as "-", then the Vserver administrator cannot
create any volumes for that Vserver. Note that the cluster administrator
will still be able to create volumes on any aggregate and assign them to
this Vserver.

``[-max-volumes <integer_or_unlimited>]`` - Limit on Maximum Number of
Volumes allowed::
This optional parameter specifies the maximum number of volumes that can
be created for the Vserver, including the root volume.

``[-admin-state
{running|stopped|starting|stopping|initializing|deleting}]`` - Vserver
Admin State [.advaced-priv]#(privilege: advanced)#::
Use this parameter to set the admin state of the Vserver if the Vserver
start or stop job fails. Possible values include `running and stopped.`

``[-allowed-protocols {nfs|cifs|fcp|iscsi|ndmp|nvme|s3}]`` - Allowed
Protocols::
This optional parameter specifies the list of protocols to be allowed to
run on the Vserver. When part of vserver-modify, this field should include
the existing list along with the new protocol list to be added to prevent
data disruptions. Possible values include nfs, cifs, fcp, iscsi, ndmp and
nvme.

``[-disallowed-protocols {nfs|cifs|fcp|iscsi|ndmp|nvme|s3}]`` - Disallowed
Protocols::
This optional parameter specifies the list of protocols to be disallowed
to run on the Vserver. When part of vserver-modify, this field should
include the existing list along with the new protocol list to be added to
```

```
prevent data disruptions. Possible values include nfs, cifs, fcp, iscsi,
ndmp and nvme.

``[-qos-policy-group <text>]`` - QoS Policy Group::
This optionally specifies which QoS policy group to apply to the Vserver.
This policy group defines measurable service level objectives (SLOs) that
apply to the storage objects with which the policy group is associated. If
you do not assign a policy group to a Vserver, the system will not monitor
and control the traffic to it. To remove this Vserver from a policy group,
enter the reserved keyword "none".

``[-caching-policy <text>]`` - Caching Policy Name::
 This optionally specifies the caching policy to apply to the Vserver. A
caching policy defines how the system caches this volume's data in Flash
Cache modules. If a caching policy is not assigned to this Vserver, the
system uses the default cluster-wide policy. The available caching
policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read, and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read, and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.

+
Default caching-policy is auto.

``[-is-space-reporting-logical {true|false}]`` - Logical Space Reporting::
This optionally specifies whether to report space logically on residing
volumes which are created after this operation. Existing volumes will not
be affected by modifying this value on an existing Vserver. To change
whether space is reported logically for existing volumes, you will have to
modify the setting on those volumes. When space is reported logically,
ONTAP reports the volume space such that all the physical space saved by
the storage efficiency features are also reported as used.

``[-is-space-enforcement-logical {true|false}]`` - Logical Space
```

```
Enforcement::
This optionally specifies whether to perform logical space accounting on
residing volumes which are created after this operation. Older volumes
will continue to have old value. When space is enforced logically, ONTAP
enforces volume settings such that all the physical space saved by the
storage efficiency features will be calculated as used.

``[-anti-ransomware-default-volume-state {disabled|dry-run}]`` - Default
Anti_ransomware State of the Vserver's Volumes::
This optional parameter specifies the default anti_ransomware state for
the volumes in the Vserver. When this is dry_run, all the volumes created
henceforth will by default have anti_ransomware state as dry-run.
Similarly, when this is disabled, all the volumes created henceforth will
by default have anti_ransomware state as disabled.

``[-auto-enable-analytics {true|false}]`` - Enable Analytics on New
Volumes::
This optional parameter specifies whether analytics is automatically
enabled for the volumes in the Vserver that is being modified. The default
value is false.

``[-auto-enable-activity-tracking {true|false}]`` - Enable Activity
Tracking on New Volumes::
This optional parameter specifies whether activity tracking is
automatically enabled for the volumes in the Vserver that is being
modified. The default value is false.

``[-storage-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Storage Limit::
This parameter specifies the storage size limit for a Vserver. This
parameter can be set to zero to disable storage-limit enforcement on a
Vserver. Default value is 0.

``[-storage-limit-threshold-alert <percent>]`` - Storage Limit Threshold
Alert::
This parameter specifies when the "vserver.storage.threshold" event log
message is generated. This parameter can be set to zero to disable
"vserver.storage.limit.threshold.exceeded" event log message. Default
value is 90%.

``[-anti-ransomware-auto-switch-from-learning-to-enabled {true|false}]`` -
Anti-ransomware Auto-switch from Learning to Enabled::
This optional parameter specifies whether anti-ransomware will be
automatically enabled on the volumes in the Vserver that is being
modified. The default value when creating a Vserver is ``_true_`` .

``[-anti-ransomware-auto-switch-minimum-incoming-data-percent <percent>]``
```

```
- Anti-ransomware Auto-switch Minimum Incoming Data (in percentage)::
This optional parameter specifies minimum incoming writes (in percentage)
required to automatically enable anti-ransomware on volumes in the
Vserver. The default value when creating a Vserver is 5%.

``[-anti-ransomware-auto-switch-duration-without-new-file-extension
<integer>]`` - Anti-ransomware Auto-switch Duration Without New File
Extension (in Days)::
This optional parameter specifies number of days without a new incoming
extension to automatically enable anti-ransomware on volumes in the
Vserver. The default value when creating a Vserver is ``_3_`` days.

``[-anti-ransomware-auto-switch-minimum-learning-period <integer>]`` -
Anti-ransomware Auto-switch Minimum Learning Period::
This optional parameter specifies minimum number of days a volume in the
Vserver must be in learning mode to automatically enable anti-ransomware.
The default value when creating a Vserver is ``_10_`` days.

``[-anti-ransomware-auto-switch-minimum-file-count <integer>]`` - Anti-
ransomware Auto-switch Minimum File Count::
This optional parameter specifies minimum number of new files created in
volume to automatically enable anti-ransomware on volume in the vserver.
The default value when creating a Vserver is ``_200_`` .

``[-anti-ransomware-auto-switch-minimum-file-extension <integer>]`` -
Anti-ransomware Auto-switch Minimum File Extension::
This optional parameter specifies minimum number of new file extensions in
a volume to automatically enable anti-ransomware on volume in the vserver.
The default value when creating a Vserver is ``_10_`` .

== Examples

The following example modifies the quota policy for a Vserver named
vs0.example.com to pol1, specifies a snapshot policy named daily, adds the
comment "Sales team access".
```

cluster1::> vserver modify -vserver vs0.example.com -snapshot-policy daily
-comment "Sales team access" -quota-policy pol1

```
[[ID7cd28274c974d9efa560a1276517feeb]]
= vserver prepare-for-revert
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Prepares Vservers to be reverted

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver prepare-for-revert` command prepares Vservers to be reverted
to the previous version of ONTAP. It disables any operations that cannot
be scheduled during revert.

== Examples

The following example prepares all Vservers to be reverted.
```

cluster1::*> vserver prepare-for-revert

```
[[IDf8ae9aca1d72633170c8bebb58dc0d52]]
= vserver remove-aggregates
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove aggregates from the Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver remove-aggregates` command removes aggregates from the
Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
Specifies the Vserver from which aggregates have to be removed.

``-aggregates <aggregate name>,...`` - List of Aggregates to Be Removed::
Specifies the list of aggregates to remove from the Vserver.

== Examples

The following example illustrates how to remove aggregates ``_aggr1_`` and
``_aggr2_`` from a Vserver named ``_vs.example.com_`` :
```

cluster1::> vserver remove-aggregates -vserver vs.example.com -aggregates aggr1,aggr2

```
[[ID2fea5eae6ab740700596c2f1dec4a595]]
= vserver remove-protocols
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove protocols from the Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver remove-protocols` command removes the specified protocols
from the specified Vserver. When you remove the protocols from a Vserver,
the data access with respect to the removed protocols is disrupted.

== Parameters

``-vserver <vserver>`` - Vserver::
Specifies the Vserver that is to be modified.

``-protocols {nfs|cifs|fcp|iscsi|ndmp|nvme|s3}`` - Protocols::
This parameter specifies the list of protocols to be removed. on the
Vserver. Possible values include nfs, cifs, fcp, iscsi, ndmp and nvme.

== Examples

The following example shows removing protocol 'cifs' from a Vserver named
vs0.example.com.
```

cluster1::> vserver remove-protocols -vserver vs0.example.com -protocols cifs

```
[[ID077b3ecce15ee5f676abe75a8f03c175]]
= vserver rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver rename` command renames the Vserver. If the vserver being
renamed is participating in an Inter-cluster Vserver peer relationship,
all the corresponding remote clusters will be updated with the new peer
Vserver name.

== Parameters

``-vserver <text>`` - Vserver::
This specifies the Vserver that is to be renamed.

``-newname <vserver>`` - New Vserver name (Use Fully Qualified Domain
Name, For example: data.example.com)::
This specifies the Vserver's new name. The name must be a unique Vserver
name in the cluster. Use a fully qualified domain name (FQDN) - for
example, "data.example.com" - for the Vserver name to reduce name
collisions in cluster leagues.
+
NOTE: Maximum number of characters supported is 47, and 41 for a Vserver
with subtype "sync-source". "all" is a reserved name and must not be used
as a Vserver name.
+


``[-foreground {true|false}]`` - Foreground Process::
This specifies whether the rename job will be run in foreground or
background. By default, the job runs in foreground.

== Examples

The following examples rename a Vserver named ``_vs1.example.com_`` as
``_vs2.example.com_`` , and then finally back to its original name:
```

```
(When there is no intercluster Vserver peer relationship with the vserver)
cluster1::> vserver rename -vserver vs1.example.com -newname vs2.example.com
(When there is at least one intercluster peer relationship with the Vserver)
cluster1::> vserver rename -vserver vs1.example.com -newname vs2.example.com
[Job 277] Job succeeded: Vserver rename completed successfully
cluster1::> vserver rename -vserver vs2.example.com -newname vs1.example.com -foreground false
[Job 278] Job is queued: Rename Vserver vs2.example.com to vs1.example.com.
```

```
[[ID84e8517d8703dd293ac9afdd7694c38a]]
= vserver restamp-msid
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restamp the MSIDs of all the volumes in a Vserver to match or be different
from the source Vserver

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver restamp-msid` command restamps MSIDs of all volumes in a dp-
destination Vserver to make them either identical to the VserverDR source
Vserver. The command is run on secondary VserverDR site and automatically
updates the MSID preserve behavior for the Vserver. A
xref:{relative_path}snapmirror-resync.html[snapmirror resync] must be run
after this command completes.

== Parameters

``-vserver <vserver name>`` - Vserver name [.advaced-priv]#(privilege:
advanced)#::
The name of the dp-destination Vserver.

``-preserve-msid {true|false}`` - Make MSID same as that of source
Vserver. False sets the values as different. [.advaced-priv]#(privilege:
advanced)#::
Boolean value through which the user can specify whether to make the MSIDs
of the volumes same as that of Source Vserver. Specifying true will make
the MSIDs same and specifying false will make them different.

== Examples

This example will stamp all the volumes of Vserver vs1dp with the same
MSID as the source Vserver.
```

cluster1::>vserver restamp-msid -vserver vs1dp -preserve-msid true

```
== Related Links
```

```
* xref:{relative_path}snapmirror-resync.html[snapmirror resync]




[[IDac290802d59cc1894d46e4db754c011b]]
= vserver show-aggregates
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show details of aggregates in a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver show-aggregates` command displays the details of all the
aggregates that are associated with Vservers. The aggregate details
displayed are the aggregate name, state, available size, the type of
aggregate and the SnapLock type.

== 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.

``[-vserver <vserver>]`` - Vserver::
If this optional parameter is specified, the command displays the details
of aggregates that are associated with the specified Vserver.

``[-aggregate <aggregate name>]`` - Aggregate::
If this optional parameter is specified, the command displays all of the
Vservers configured with the specified aggregate.

== Examples

The following example displays the aggregates configured for Vserver vs.
```

```
cluster1::> vserver show-aggregates -vserver vs
Available
Vserver Aggregate State Size Type SnapLock-Type
--------------- --------------- ------- ---------- ------- --------------
vs aggr1 online 795.2MB hdd non-snaplock
vs aggr2 online 795.2MB hdd non-snaplock
2 entries were displayed.
```

```
[[ID81792041a06bed3b3a7aa9b511e19095]]
= vserver show-protocols
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show protocols for Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver show-protocols` command displays the running protocols on a
given Vserver.

== 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.

``[-vserver <vserver>]`` - Vserver::
If this parameter is specified, the command displays the allowed set of
protocols for the specified Vserver.

``[-protocol {nfs|cifs|fcp|iscsi|ndmp|nvme|s3}]`` - Protocols::
If this optional parameter is specified, the command displays all the
Vservers configured with the specified protocols.

== Examples

The following example displays the protocols configured for Vserver vs1.
```

cluster1::> vserver show-protocols -vserver vs1
Vserver: vs1
Protocols: nfs, cifs

```
[[ID3a0fa959bccfcc7e3a1660a667895019]]
= vserver show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Vservers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver show` command displays the following information:

* Vserver name
* Vserver type (``_data_`` , ``_admin_`` , ``_node_`` or ``_system_`` -
detailed view only)
* Vserver subtype (``_default_`` , ``_dp-destination_`` , ``_sync-
source_`` , and ``_sync-destination_`` - detailed view only)
* Vserver universal unique identifier (detailed view only)
* Root volume name
* Aggregate on which the root volume is located
* Associated NIS domain
* Root volume security style (``_unix_`` for UNIX mode bits, ``_ntfs_``
for CIFS ACLs, ``_mixed_`` for both (detailed view only)
* LDAP client
* Language (detailed view only)
* Snapshot policy (detailed view only)
* Comment text (detailed view only)
* Quota policy (detailed view only)
* Aggregate list (detailed view only)
* Maximum Volumes (detailed view only)
* Qos-policy-group (detailed view only)
* Config-lock (detailed view only)
* Admin state (``_running_`` , ``_stopped_`` , ``_starting_`` ,
``_stopping_`` , ``_initializing_`` , or ``_deleting_`` )
* Operational state (``_running_`` , or ``_stopped_`` )
* Operational state stopped reason (``_sync-destination-and-switchover-
not-done_`` , or ``_cluster-reboot-done_`` , or ``_admin-state-stopped_``
)
* Allowed Protocols (``_nfs_`` , ``_cifs_`` , ``_fcp_`` , ``_iscsi_`` ,
``_nvme_`` , ``_ndmp_`` - detailed view only)
* Disallowed Protocols (``_nfs_`` , ``_cifs_`` , ``_fcp_`` , ``_iscsi_`` ,
```

```
``_nvme_`` , ``_ndmp_`` - detailed view only)
* IPspace to which the Vserver belongs (detailed view only)
* Caching policy



== 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.

| ``[-protocols  ]``::
If this optional parameter is specified, the command displays the allowed
and disallowed set of protocols for the Vserver(s).

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver>]`` - Vserver::
If this parameter is specified, the command displays detailed information
about the specified Vserver.

``[-type <vserver type>]`` - Vserver Type::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that have the specified Vserver type. Types
include ``_admin_`` for the cluster-wide management Vserver, ``_system_``
for cluster-level communications in an IPspace, ``_data_`` for data
serving Vserver, and ``_node_`` for node management Vserver.

``[-subtype <vserver subtype>]`` - Vserver Subtype::
 If this parameter is specified, the command displays information only
about the Vserver or Vservers that have the specified Vserver subtype.
Types include:
+
  * ``_default_`` for default data Vserver
  *
              and

  * ``_dp-destination_`` for Data Protection destination Vserver.
  * ``_sync-source_`` for MetroCluster source Vserver,
  * ``_sync-destination_`` for MetroCluster destination Vserver.
```

```
``[-uuid <UUID>]`` - Vserver UUID::
If this parameter is specified, the command displays information only
about the Vserver that match the specified UUID.

``[-rootvolume <volume name>]`` - Root Volume::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that have the specified root volume.

``[-aggregate <aggregate name>]`` - Aggregate::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that have their root volumes contained by
the specified aggregate.

``[-nisdomain <nis domain>]`` - NIS Domain::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that use the specified NIS domain.

``[-rootvolume-security-style <security style>]`` - Root Volume Security
Style::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that have the specified root-volume security
style.

``[-ldap-client <text>]`` - LDAP Client::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that use the specified LDAP client.

``[-language <Language code>]`` - Default Volume Language Code::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that use the specified language. To
determine the available languages, enter "`vserver show``-language```_?_``
" at the clustershell command prompt and at the Vserver prompt.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that have the specified snapshot policy.

``[-data-services <LIF Service Name>,...]`` - Data Services::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified data services.

``[-comment <text>]`` - Comment::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified comment.

``[-quota-policy <text>]`` - Quota Policy::
If this parameter is specified, the command displays information only
```

```
about the Vserver or Vservers that use the specified quota policy.

``[-aggr-list <aggregate name>,...]`` - List of Aggregates Assigned::
If this parameter is specified, the command displays information only
about the Vserver or Vservers to which the specified aggregate(s) are
assigned for use.

``[-max-volumes <integer_or_unlimited>]`` - Limit on Maximum Number of
Volumes allowed::
If this parameter is specified, the command displays information only
about the Vserver or Vservers on which the specified maximum volume count
is configured.

``[-admin-state
{running|stopped|starting|stopping|initializing|deleting}]`` - Vserver
Admin State::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified admin-state.

``[-operational-state {running|stopped}]`` - Vserver Operational State::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified operational-state.
This field determines the state of the Vserver LIFs. New LIFs created on a
Vserver, which is in running state, will be operationally up and the LIFs
created on a Vserver, which is in stopped state, will be operationally
down.

``[-operational-state-stopped-reason {sync destination and switchover is
not done|cluster reboot is done|admin state stopped| dp destination not
started}]`` - Vserver Operational State Stopped Reason::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that are operationally stopped due to the
specified reason. This field indicates the reason for the operational-
state of the Vserver being stopped

``[-allowed-protocols {nfs|cifs|fcp|iscsi|ndmp|nvme|s3}]`` - Allowed
Protocols::
If this parameter is specified, the command displays information only
about the Vserver or Vservers on which the specified protocols are allowed
to run.

``[-disallowed-protocols {nfs|cifs|fcp|iscsi|ndmp|nvme|s3}]`` - Disallowed
Protocols::
If this parameter is specified, the command displays information only
about the Vserver or Vservers on which the specified protocols are
disallowed to run.
```

```
``[-is-repository {true|false}]`` - Is Vserver with Infinite Volume::
If this parameter is no longer supported.

``[-qos-policy-group <text>]`` - QoS Policy Group::
Display the Vservers that match the specified qos-policy-group.
+
A policy group defines measurable service level objectives (SLOs) that
apply to the storage objects with which the policy group is associated. If
you do not assign a policy group to a Vserver, the system will not monitor
and control the traffic to it.

``[-caching-policy <text>]`` - Caching Policy Name::
Display the Vservers that match the specified caching-policy.
+
A caching policy defines the caching behavior of this Vserver at the Flash
Cache level. If a caching policy is not assigned to this Vserver, the
system uses the default cluster-wide policy. The available caching
policies are:
+
  * none - Does not cache any user data or metadata blocks.
  * auto - Read caches all metadata and randomly read user data blocks,
and write caches all randomly overwritten user data blocks.
  * meta - Read caches only metadata blocks.
  * random_read - Read caches all metadata and randomly read user data
blocks.
  * random_read_write - Read caches all metadata, randomly read, and
randomly written user data blocks.
  * all_read - Read caches all metadata, randomly read, and sequentially
read user data blocks.
  * all_read_random_write - Read caches all metadata, randomly read,
sequentially read, and randomly written user data.
  * all - Read caches all data blocks read and written. It does not do any
write caching.

+
Default caching-policy is auto.

``[-config-lock {true|false}]`` - Config Lock::
This parameter specifies if the Vserver is locked or unlocked for
modification. If the config-lock is set to true, then modifying the
Vserver's configuration is not allowed.

``[-ipspace <IPspace>]`` - IPspace Name::
If this parameter is specified, the command displays information only
about the Vservers that are assigned to the specified IPspace.
```

```
``[-foreground {true|false}]`` - Foreground Process::
This optionally specifies whether the Vserver show operation can be
executed in the background. If nothing is specified, by default the
Vserver show operation is executed in the foreground.

``[-is-space-reporting-logical {true|false}]`` - Logical Space Reporting::
This optionally specifies whether to report space logically on residing
volumes. When space is reported logically, ONTAP reports the volume space
such that all the physical space saved by the storage efficiency features
are also as reported as used.

``[-is-space-enforcement-logical {true|false}]`` - Logical Space
Enforcement::
This optionally specifies whether to perform logical space accounting on
residing volumes. When space is enforced logically, ONTAP enforces volume
settings such that all the physical space saved by the storage efficiency
features will be calculated as used.

``[-anti-ransomware-default-volume-state {disabled|dry-run}]`` - Default
Anti_ransomware State of the Vserver's Volumes::
This optionally specifies the default anti_ransomware state for the
volumes in the Vserver. The available anti_ransomware states for the
volumes in the Vserver are:
+
  * disabled - Anti_ransomware protection is not enabled on the volumes in
the Vserver.
  * dry-run - Anti_ransomware protection is in evaluation mode on the
volumes in the Vserver.



``[-auto-enable-analytics {true|false}]`` - Enable Analytics on New
Volumes::
This parameter specifies whether analytics is automatically enabled for
the volumes in the Vserver.

``[-auto-enable-activity-tracking {true|false}]`` - Enable Activity
Tracking on New Volumes::
This parameter specifies whether activity tracking is automatically
enabled for the volumes in the Vserver.

``[-storage-allocated {<integer>[KB|MB|GB|TB|PB]}]`` - Total Size of the
Volumes::
If this parameter is specified, the command displays the information only
about the Vserver or Vservers that match the specified value for storage
allocated.
```

```
``[-storage-available {<integer>[KB|MB|GB|TB|PB]}]`` - Available Size::
If this parameter is specified, the command displays the information only
about the Vserver or Vservers that match the specified value for storage
available.

``[-storage-used-percentage <percent>]`` - Used Percent::
If this parameter is specified, the command displays the information only
about the Vserver or Vservers that match the specified value for used
storage percentage.

``[-number-of-volumes-in-recovery-queue <integer>]`` - Number of Volumes
in Recovery Queue::
If this parameter is specified, the command displays the information only
about the Vserver or Vservers that match the specified value for number of
volumes in recovery queue.

``[-total-volume-size-in-recovery-queue {<integer>[KB|MB|GB|TB|PB]}]`` -
Storage Space in Recovery Queue Volumes::
If this parameter is specified, the command displays the information only
about the Vserver or Vservers that match the specified value for total
volume size in recovery queue.

``[-storage-limit-threshold-exceeded {true|false}]`` - Max Storage Alert
Threshold Exceeded::
If this parameter is specified, the command displays the information only
about the Vserver or Vservers whether that exceeded the value specified in
the storage-limit-threshold-alert.

``[-storage-limit {<integer>[KB|MB|GB|TB|PB]}]`` - Storage Limit::
If this parameter is specified, the command displays the information only
about the Vserver or Vservers that match the specified value for storage
limit.

``[-storage-limit-threshold-alert <percent>]`` - Storage Limit Threshold
Alert::
If this parameter is specified, the command display the information only
about the Vserver or Vservers that match the specified value for storage
limit threshold alert.

``[-anti-ransomware-auto-switch-from-learning-to-enabled {true|false}]`` -
Anti-ransomware Auto-switch from Learning to Enabled::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified value for
anti_ransomware_auto_switch_from_learning_to_enabled.

``[-anti-ransomware-auto-switch-minimum-incoming-data-percent <percent>]``
```

```
- Anti-ransomware Auto-switch Minimum Incoming Data (in percentage)::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified value for
anti_ransomware_auto_switch_minimum_incoming_data.

``[-anti-ransomware-auto-switch-duration-without-new-file-extension
<integer>]`` - Anti-ransomware Auto-switch Duration Without New File
Extension (in Days)::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified value for
anti_ransomware_auto_switch_duration_without_new_file_extension.

``[-anti-ransomware-auto-switch-minimum-learning-period <integer>]`` -
Anti-ransomware Auto-switch Minimum Learning Period::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified value for
anti_ransomware_auto_switch_minimum_learning_period.

``[-anti-ransomware-auto-switch-minimum-file-count <integer>]`` - Anti-
ransomware Auto-switch Minimum File Count::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified value for
anti_ransomware_auto_switch_minimum_file_count.

``[-anti-ransomware-auto-switch-minimum-file-extension <integer>]`` -
Anti-ransomware Auto-switch Minimum File Extension::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified value for
anti_ransomware_auto_switch_minimum_file_extension.

``[-storage-reserved {<integer>[KB|MB|GB|TB|PB]}]`` - Storage Reserved for
In-flight Volume Operations::
If this parameter is specified, the command displays information only
about the Vserver or Vservers that match the specified value for storage
reserved.

== Examples

The following example displays information about all Vservers.
```

cluster1::> vserver show

non mcc setup:
Admin Operational Root
Vserver Type Subtype state state Volume Aggregate
---------- ------- --------- ---------- ---------- ---------- ---------
cluster admin - - - - -

```
node1 node - - - - -
vs0 data default running running root_vs1 aggr0
vs1 data dp-destination stopped stopped - -

4 entries were displayed.
```

mcc setup:

```
cluster1::> vserver show
Admin Operational Root
Vserver Type Subtype state state Volume Aggregate
---------- ------- --------------- ------- ----------- ----------- ---------
cluster admin - - - - -
node1 node - - - - -
vs2 data sync-source running running rv data_aggr
vs3-mc data sync-destination running stopped - -

4 entries were displayed.
```

```
[[IDa03247fde25cc81df7715ab1c4b0da9e]]
= vserver start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver start` command starts data access on a Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
This specifies the name of the Vserver on which the data access is to be
started. This operation is only supported on a data Vserver.
+
NOTE: The name must be of 47 characters length or less.
+


``[-foreground {true|false}]`` - Foreground Process::
This specifies if the `vserver start` command should be executed in the
foreground or background. If you do not enter this parameter, it is set to
``_true_`` , and the `vserver start` command is executed in the
foreground.

``[-force <true>]`` - Force Vserver Start::
In case of a MetroCluster configuration or Vserver disaster recovery, by
using this parameter you can start the Vserver that is either locked
(which prevents any configuration changes) or its partner Vserver is
operationally running. If you do not enter this parameter, it is set to
false.

== Examples

The following example starts data access on Vserver ``_vs0.example.com_``
in the background.
```

cluster1::> vserver start -vserver vs0.example.com -foreground false

```
[[ID279a0166cca248e058d346c8b20ca604]]
= vserver stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop a Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver stop` command stops data access on a Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
This specifies the name of the Vserver on which the data access is to be
stopped. This operation is only supported on a data Vserver.
+
NOTE: The name must be of 47 characters length or less.
+


``[-foreground {true|false}]`` - Foreground Process::
This specifies if `vserver stop` command should be executed in the
foreground or background. If you do not enter this parameter, it is set to
``_true_`` , and the ` vserver stop` command is executed in the
foreground.

== Examples

The following example stops data access on Vserver ``_vs0.example.com_``
in the background.
```

cluster1::> vserver stop -vserver vs0.example.com -foreground false

```
[[ID22abccc6c41ce050907dd3e48a2b7ca8]]
= vserver unlock
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Unlock Vserver configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver unlock` command revokes the administrative lock on the
Vserver configuration. When a Vserver is unlocked, changes to the
configuration are permitted. The unlock operation fails if the Vserver is
not locked by the administrator or if it is locked by internal
applications. If the Vserver fails to unlock due to an error condition,
you can use the -force option.

== Parameters

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of the Vserver that has to be unlocked.

``[-force <true>]`` - Force Unlock [.advaced-priv]#(privilege:
advanced)#::
This option is specified to unlock the Vserver when the Vserver fails to
unlock due to an error condition.

== Examples

The following example illustrates how to unlock the Vserver named
``_vs123.example.com, forcefully_`` :
```

cluster1::> vserver unlock -vserver vs1.example.com -force true

```
= vserver active-directory commands

:leveloffset: +1


[[ID107fce649427a91379f44070a8417654]]
```

```
= vserver active-directory create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an Active Directory account. If joining a domain, this command may
take several minutes to complete.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver active-directory create` command creates an Active Directory
account for a Vserver. When you create the Active Directory account, you
must add it to an existing Windows Active Directory domain. When you enter
this command, you are prompted to provide the credentials of a user
account that has sufficient privileges to add computers to the `-ou`
container within the `-domain` domain. The user account must have a
password that cannot be empty. When joining a domain, this command may
take several minutes to complete.

NOTE: Each Vserver can have only one Active Directory account.



== Parameters

``-vserver <vserver>`` - Vserver::
This parameter specifies the name of the Vserver for which you want to
create the Active Directory account. The Vserver must already exist.

``-account-name <NetBIOS>`` - Active Directory NetBIOS Name::
This parameter specifies the name of the Active Directory account (up to
15 characters).

``-domain <TextNoCase>`` - Fully Qualified Domain Name::
This parameter specifies the name of the Active Directory domain.

``[-ou <text>]`` - Organizational Unit::
This parameter specifies the organizational unit within the Active
Directory domain. By default, this parameter is set to ``_CN=Computers_``
. When specifying this parameter, specify only the organizational unit
portion of the distinguished name. ONTAP appends the value provided for
the required `-domain` parameter onto the value provided for `-ou`
```

```
parameter to produce the Active Directory distinguished name, which is
used when creating the Vserver’s Active Directory account in the domain.
+
NOTE: Nested OUs must be provided in a specific order with all containers
separated by a comma. Reading from left to right you travel up the
directory tree until you reach the root OU.
+


== Examples

The following example creates an Active Directory account ``_ADSERVER1_``
for Vserver ``_vs1_`` and domain ``_example.com_`` .
```

cluster1::> vserver active-directory create -vserver vs1 -account-name ADSERVER1 -domain example.com

In order to create an Active Directory machine account, you must supply the
name and password of a Windows account with sufficient privileges to add
computers to the "CN=Computers" container within the "example.com" domain.

Enter the user name: Administrator

Enter the password:

```
The following example creates an Active Directory account ``_ADSERVER2_``
for Vserver ``_vs2_`` , domain ``_example.com_`` and organizational unit
``_sample_ou_`` .
```

cluster1::> vserver active-directory create -vserver vs2 -account-name ADSERVER2 -domain example.com -ou OU=sample_ou

In order to create an Active Directory machine account, you must supply the
name and password of a Windows account with sufficient privileges to add
computers to the "OU=sample_ou" container within the "example.com" domain.

Enter the user name: Administrator

Enter the password:

```
The following example creates an Active Directory account ``_ADSERVER2_``
for Vserver ``_vs2_`` , domain ``_example.com_`` and nested organizational
unit ``_OU=developers,OU=engineering,OU=corp_`` .
```

cluster1::> vserver active-directory create -vserver vs2 -account-name ADSERVER2 -domain example.com -ou OU=developers,OU=engineering,OU=corp

In order to create an Active Directory machine account, you must supply the

name and password of a Windows account with sufficient privileges to add
computers to the "OU=developers,OU=engineering,OU=corp" container within the "example.com" domain.

Enter the user name: Administrator

Enter the password:

```
[[ID62e1fd1e075fc15e9be682c3ef3a9ced]]
= vserver active-directory delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an Active Directory account

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver active-directory delete` command deletes the Active Directory
account for a specified Vserver.

== Parameters

``-vserver <vserver>`` - Vserver::
This parameter specifies the Vserver for the Active Directory account you
want to delete.

== Examples

The following example deletes the Active Directory account for a Vserver
named ``_vs1_`` :
```

cluster1::> vserver active-directory delete -vserver vs1
In order to delete an Active Directory machine account, you must supply the
name and password of a Windows account with sufficient privileges to remove
computers from the "example.com" domain.

Enter the user name: Administrator

Enter the password:

```
[[ID41ffd85da84f845ff085c9ce59621735]]
= vserver active-directory modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the domain of an Active Directory account. If re-joining the
current domain or joining a new one, this command may take several minutes
to complete.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver active-directory modify` command modifies the domain of an
Active Directory account. You can also re-join the current domain or join
a new one. When joining a domain, this command may take several minutes to
complete.

== Parameters

``-vserver <vserver>`` - Vserver::
This parameter specifies the Vserver for the Active Directory account
whose associated domain you want to modify.

``[-domain <TextNoCase>]`` - Fully Qualified Domain Name::
This parameter specifies the fully qualified name of the Active Directory
domain to associate with the Active Directory account.

== Examples

The following example modifies the Active Directory domain associated with
Vserver ``_vs1_`` .
```

cluster1::> vserver active-directory modify -vserver vs1 -domain example.com

In order to create an Active Directory machine account, you must supply the
name and password of a Windows account with sufficient privileges to add
computers to the "CN=Computers" container within the "example.com" domain.

Enter the user name: administrator

Enter the password:

```
[[ID0dfb1cbb1aa68e52fa07d45acc16e310]]
= vserver active-directory password-change
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change the domain account password for an Active Directory account

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver active-directory password-change` command changes the domain
account password for the specified Vserver's Active Directory account.

== Parameters

``-vserver <vserver>`` - Vserver::
This parameter specifies the name of the Vserver associated with the
Active Directory account whose domain account password you want to change.

== Examples

The following example changes the password for the Active Directory
account for a Vserver named ``_vs1_`` .
```

cluster1::> vserver active-directory password-change -vserver vs1

```
[[IDfa344c062ddee445fb173f1aafeae096]]
= vserver active-directory password-reset
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset the domain account password for an Active Directory account

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver active-directory password-reset` command resets the domain
account password for the Active Directory account. This may be required if
the password stored along with the machine account in the Windows Active
Directory domain is changed or reset without the Vserver's knowledge. The
operation requires the credentials for a user with permission to reset the
password in the organizational unit (OU) that contains the machine
account.

== Parameters

``-vserver <vserver>`` - Vserver::
This parameter specifies the name of the Vserver associated with the
Active Directory account whose domain account password you want to reset.

== Examples

The following example resets the password for the Active Directory account
for a Vserver named ``_vs1_`` .
```

cluster1::> vserver active-directory password-reset -vserver vs1

Enter your user ID: Administrator
Enter your password:

```
[[ID77bebaa14189efc1f007e2dad7c8122a]]
= vserver active-directory show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display Active Directory accounts

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver active-directory show` command displays information about
Active Directory accounts. The command output depends on the parameter or
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all
Active Directory accounts:

* Vserver name
* Active Directory account NetBIOS name
* Domain or workgroup name



You can specify the `-fields` parameter to specify which fields of
information to display about Active Directory accounts. You can use
`-fields```_?_`` to display the valid values for the `-fields` parameter.
In addition to the fields above, you can display the following fields:

* Fully-qualified domain name
* Organizational unit

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
Active Directory accounts that are in the Windows Active Directory domain
named ``_RUBY_`` , run the command with the value of the `-domain-
workgroup` parameter set to ``_RUBY_`` .

You can specify the `-instance` parameter to display all information for
all Active Directory accounts in list form.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields` <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.
```

```
``[-vserver <vserver>]`` - Vserver::
If you specify this parameter, the command displays information only about
the Active Directory account for the specified Vserver.

``[-account-name <NetBIOS>]`` - Active Directory NetBIOS Name::
If you specify this parameter, the command displays information only for
the Active Directory accounts that match the specified NetBIOS account
name.

``[-domain-workgroup <CIFS domain>]`` - NetBIOS Domain/Workgroup Name::
If you specify this parameter, the command displays information only for
the Active Directory accounts that are in the specified NetBIOS domain or
workgroup.
+
NOTE: Workgroups are not supported in this release.
+


``[-domain <TextNoCase>]`` - Fully Qualified Domain Name::
If you specify this parameter, the command displays information only for
the Active Directory accounts that are in the specified domain.

``[-ou <text>]`` - Organizational Unit::
If you specify this parameter, the command displays information only for
the Active Directory accounts that are in the specified organizational
unit.

``[-auth-style {domain|workgroup|realm}]`` - Authentication Style::
If you specify this parameter, the command displays information only for
the Active Directory accounts that are in the specified authentication
style.

== Examples

The following example displays a subset of the information about all
Active Directory accounts.
```

cluster1::> vserver active-directory show
Account Domain/Workgroup
Vserver Name Name
--- --- ---
vs1 ADSERVER1 EXAMPLE

```
The following example displays all information about all Active Directory
Vservers in list form.
```

cluster1::> vserver active-directory show -instance
Vserver: vs1
Active Directory account NetBIOS Name: ADSERVER1
NetBIOS Domain/Workgroup Name: EXAMPLE
Fully Qualified Domain Name: EXAMPLE.COM
Organizational Unit: CN=Computers

```
:leveloffset: -1


= vserver audit commands

:leveloffset: +1


[[IDa8c0b9b04870df702a6fab6e9fa579c1]]
= vserver audit create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver audit create` command creates an audit configuration for a
Vserver.

When you create an audit configuration, you can also specify the rotation
method. By default, the audit log is rotated based on size.

You can use the time-based rotation parameters in any combination (`-
rotate-schedule-month` , `-rotate-schedule-dayofweek` , `-rotate-schedule-
day` , `-rotate-schedule-hour` , and `-rotate-schedule-minute` ). The `-
rotate-schedule-minute` parameter is mandatory. All other time-based
rotation parameters are optional.

The rotation schedule is calculated by using all the time-related values.
```

```
For example, if you specify only the `-rotate-schedule-minute` parameter,
the audit log files are rotated based on the minutes specified on all days
of the week, during all hours on all months of the year. If you specify
only one or two time-based rotation parameters (say `-rotate-schedule-
month` and `-rotate-schedule-minutes` ), the log files are rotated based
on the minute values that you specified on all days of the week, during
all hours, but only during the specified months. For example, you can
specify that the audit log is to be rotated during the months January,
March, and August on all Mondays, Wednesdays, and Saturdays at 10:30.

If you specify values for both `-rotate-schedule-dayofweek` and `-rotate-
schedule-day` , they are considered independently. For example if you
specify `-rotate-schedule-dayofweek` as Friday and `-rotate-schedule-day`
as 13 then the audit logs would be rotated on every Friday and on the 13th
day of the specified month, not just on every Friday the 13th.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which to create the
audit configuration. The Vserver must already exist.

``-destination <text>`` - Log Destination Path::
This parameter specifies the audit log destination path where consolidated
audit logs are stored. If the path is not valid, the command fails. The
path can be up to 864 characters in length and must have read-write
permissions.

``[-events {file-ops|cifs-logon-logoff|cap-staging|file-share|audit-
policy-change|user-account|authorization-policy-change|security-
group|async-delete}]`` - Categories of Events to Audit::
This parameter specifies the categories of events to be audited. Supported
event categories are: file access events (both CIFS and NFS), CIFS logon
and logoff events, Central Access Policy(CAP) staging events, File share
events, Audit policy change events, Local User Account Management Events,
Local Security Group Management Events and Authorization Policy Change
Events. The corresponding parameter values are: ``_file-ops_`` , ``_cifs-
logon-logoff_`` , ``_cap-staging_`` , ``_file-share_`` , ``_audit-policy-
change_`` , ``_user-account_`` , ``_security-group_`` and
``_authorization-policy-change_`` .By default, ``_file-ops_`` , ``_cifs-
logon-logoff_`` and ``_audit-policy-change_`` events are enabled. The
support for ``_audit-policy-change_`` event can be modified from diag
prompt using xref:{relative_path}vserver-audit-modify.html[vserver audit
modify] command.

``[-format {xml|evtx}]`` - Log Format::
```

```
This parameter specifies the output format of the audit logs. The output
format can be either Data ONTAP-specific XML or Microsoft Windows EVTX log
format. By default, the output format is EVTX.

``[-rotate-size {<size>|-}]`` - Log File Size Limit::
This parameter specifies the audit log file size limit. By default, the
audit log is rotated based on size. The default audit log size is 100 MB.

``[-rotate-schedule-month <cron_month>,...]`` - Log Rotation Schedule:
Month::
This parameter specifies the monthly schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated during
the months January, March, and August, or during all the months. Valid
values are January, February, March, April, May, June, July, August,
September, October, November, December, and all. Specify "all" to rotate
the audit logs every month.

``[-rotate-schedule-dayofweek <cron_dayofweek>,...]`` - Log Rotation
Schedule: Day of Week::
This parameter specifies the daily (day of the week) schedule for rotating
the audit log. For example, you can specify that the audit log is to be
rotated on Tuesdays and Fridays, or during all the days of a week. Valid
values are Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday,
and all. Specify "all" to rotate the audit logs every day.

``[-rotate-schedule-day <cron_dayofmonth>,...]`` - Log Rotation Schedule:
Day::
This parameter specifies the day of the month schedule for rotating the
audit log. For example, you can specify that the audit log is to be
rotated on the 10th and 20th days of a month, or all days of a month.
Valid values range from 1 to 31.

``[-rotate-schedule-hour <cron_hour>,...]`` - Log Rotation Schedule:
Hour::
This parameter specifies the hourly schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at 6 a.m
and 10 a.m. Valid values range from 0 (midnight) to 23 (11:00 p.m.).
Specify "all" to rotate the audit logs every hour.

``[-rotate-schedule-minute <cron_minute>,...]`` - Log Rotation Schedule:
Minute::
This parameter specifies the minute schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at the
30th minute. Valid values range from 0 to 59.

{ ``[-rotate-limit <integer>]`` - Log Files Rotation Limit::
```

```
This parameter specifies the audit log files rotation limit. A value of 0
indicates that all the log files are retained. The default value is 10 for
cloud optimized platform and 0 for all other platform. For example, if you
enter a value of 5, the last five audit logs are retained.

| ``[-retention-duration
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Log Retention
Duration }::
This parameter specifies the audit log files retention duration. A value
of 0s indicates that all the log files are retained. The default value is
0s. For example, if you enter a value of 5d0h0m, logs more than 5 days old
are deleted.

== Examples

The following examples create an audit configuration for Vserver vs1 using
size-based rotation.
```

cluster1::> vserver audit create -vserver vs1 -destination /audit_log -rotate-size 10MB -rotate-limit 5

```
 + +

The following example creates an audit configuration for Vserver vs1 using
time-based rotation. The audit logs are rotated monthly, all days of the
week, at 12:30.
```

cluster1::> vserver audit create -vserver vs1 -destination /audit_log -rotate-schedule-month all -rotate-schedule -dayofweek all -rotate-schedule-hour 12 -rotate-schedule-minute 30

```
The following example creates an audit configuration for Vserver vs1 using
time-based rotation. The audit logs are rotated in January, March, May,
July, September, and November on Monday, Wednesday, and Friday, at 6:15,
6:30, 6:45, 12:15, 12:30, 12:45, 18:15, 18:30, and 18:45. The last 6 audit
logs are retained.
```

cluster1::> vserver audit create -vserver vs1 -destination /audit_log -rotate-schedule-month January,March,May,July,September,November -rotate-schedule-dayofweek Monday,Wednesday,Friday -rotate -schedule-hour 6,12,18 -rotate-schedule-minute 15,30,45 -rotate-limit 6

```
The following example creates an audit configuration for Vserver vs1 for
auditing CIFS and NFS file access events in the output log format EVTX.
```

cluster1::> vserver audit create -vserver vs1 -destination /audit_log -format evtx -events file-ops

```
== Related Links

* xref:{relative_path}vserver-audit-modify.html[vserver audit modify]




[[ID46a059162c08fbbbf1d048a182c86b6e]]
= vserver audit delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver audit delete` command deletes the audit configuration for a
Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver associated with the audit
configuration to be deleted.

``[-force <true>]`` - Force Delete [.advaced-priv]#(privilege:
advanced)#::
This parameter is used to forcibly delete the audit configuration. By
default the setting is `false` .

== Examples

The following example deletes the audit configuration for Vserver vs1.
```

cluster1::> vserver audit delete -vserver vs1

```
[[ID0ecf8d14596299fc564b8811137f5eb3]]
= vserver audit disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable auditing

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver audit disable` command disables auditing for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which auditing is to
be disabled. The Vserver audit configuration must already exist.

== Examples

The following example disables auditing for Vserver vs1.
```

cluster1::> vserver audit disable -vserver vs1

```
[[IDafb9c44ba95bfc10c6b1943287c9aeb2]]
= vserver audit enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable auditing

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver audit enable` command enables auditing for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which auditing is to
be enabled. The Vserver audit configuration must already exist.

``[-force <true>]`` - Force Enable [.advaced-priv]#(privilege:
advanced)#::
This parameter is used to ignore errors while enabling auditing.

== Examples

The following example enables auditing for Vserver vs1:
```

cluster1::> vserver audit enable -vserver vs1

```
[[IDd125a310e3d23b9002858551b7f41172]]
= vserver audit modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver audit modify` command modifies an audit configuration for a
Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which the audit
configuration is to be modified. The Vserver audit configuration must
already exist.
+
If you have configured time-based rotation, modifying one parameter of
time-based rotation schedule does not affect the other parameters. For
example, if the rotation schedule is set to run at Monday 12:30 a.m., and
you modify the `-rotate-schedule-dayofweek` parameter to
Monday,Wednesday,Friday, the new rotation-schedule rotates the audit logs
on Monday, Wednesday, and Friday at 12:30 a.m. To clear time-based
rotation parameters, you must explicitly set that portion to "-". Some
time-based parameters can also be set to "all".

``[-destination <text>]`` - Log Destination Path::
This parameter specifies the audit log destination path where consolidated
audit logs are stored. If the path is not valid, the command fails. The
path can be up to 864 characters in length and must have read-write
permissions.

``[-events {file-ops|cifs-logon-logoff|cap-staging|file-share|audit-
policy-change|user-account|authorization-policy-change|security-
group|async-delete}]`` - Categories of Events to Audit::
This parameter specifies the categories of events to be audited. Supported
event categories are: file access events (both CIFS and NFS), CIFS logon
and logoff events, Central Access Policy(CAP) staging events, File share
events, Audit policy change events, Local User Account Management Events,
Local Security Group Management Events and Authorization Policy Change
Events. The corresponding parameter values are: ``_file-ops_`` , ``_cifs-
logon-logoff_`` , ``_cap-staging_`` , ``_file-share_`` , ``_audit-policy-
change_`` , ``_user-account_`` , ``_security-group_`` and
``_authorization-policy-change_`` . By default, ``_file-ops_`` , ``_cifs-
logon-logoff_`` and ``_audit-policy-change_`` events are enabled

``[-format {xml|evtx}]`` - Log Format::
This parameter specifies the output format of the audit logs. The output
format can be either Data ONTAP-specific XML or Microsoft Windows EVTX log
format. By default, the output format is EVTX.
```

```
``[-rotate-size {<size>|-}]`` - Log File Size Limit::
This parameter specifies the audit log file size limit. By default, the
audit log is rotated based on size. The default audit log size is 100 MB.

``[-rotate-schedule-month <cron_month>,...]`` - Log Rotation Schedule:
Month::
This parameter specifies the monthly schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated during
the months January, March, and August, or during all the months. Valid
values are January, February, March, April, May, June, July, August,
September, October, November, December, and all. Specify "all" to rotate
the audit logs every month.

``[-rotate-schedule-dayofweek <cron_dayofweek>,...]`` - Log Rotation
Schedule: Day of Week::
This parameter specifies the daily (day of the week) schedule for rotating
the audit log. For example, you can specify that the audit log is to be
rotated on Tuesdays and Fridays, or during all the days of a week. Valid
values are Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday,
and all. Specify "all" to rotate the audit logs every day.

``[-rotate-schedule-day <cron_dayofmonth>,...]`` - Log Rotation Schedule:
Day::
This parameter specifies the day of the month schedule for rotating the
audit log. For example, you can specify that the audit log is to be
rotated on the 10th and 20th days of a month, or all days of a month.
Valid values range from 1 to 31.

``[-rotate-schedule-hour <cron_hour>,...]`` - Log Rotation Schedule:
Hour::
This parameter specifies the hourly schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at 6 a.m
and 10 a.m. Valid values range from 0 (midnight) to 23 (11:00 p.m.).
Specify "all" to rotate the audit logs every hour.

``[-rotate-schedule-minute <cron_minute>,...]`` - Log Rotation Schedule:
Minute::
This parameter specifies the minute schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at the
30th minute. Valid values range from 0 to 59.

{ ``[-rotate-limit <integer>]`` - Log Files Rotation Limit::
This parameter specifies the audit log files rotation limit. A value of 0
indicates that all the log files are retained. The default value is 0.

| ``[-retention-duration
```

```
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Log Retention
Duration }::
This parameter specifies the audit log files retention duration. A value
of 0s indicates that all the log files are retained. For example, if you
enter a value of 5d0h0m0s, logs more than 5 days old are deleted.

``[-audit-guarantee {true|false}]`` - Strict Guarantee of Auditing::
This parameter specifies strict guarantee of auditing for a Vserver. If
this value is true, file access is denied if audit records cannot be
generated. If this value is false, auditing is done on a best-effort
basis.

``[-charge-qos {true|false}]`` - Audit Log Extra Charge::
This parameter specifies whether the audit logs will incur an extra
charge. This extra charge is counted against the Volume's Qos policy. If
this value is set to ``_false_`` , it indicates that audit logs will not
incur an extra charge.

== Examples

The following example modifies the rotate-size and rotate-limit field for
Vserver vs1.
```

cluster1::> vserver audit modify -vserver vs1 -rotate-size 10MB -rotate-limit 3

```
The following example modifies an audit configuration for Vserver vs1
using the time-based rotation method. The audit logs are rotated monthly,
all days of the week, at 12:30.
```

cluster1::> vserver audit modify -vserver vs1 -destination /audit_log -rotate-schedule-month all -rotate-schedule -dayofweek all -rotate-schedule-hour 12 -rotate-schedule-minute 30

```
The following example modifies an audit configuration for Vserver vs1 for
auditing CIFS and NFS file access events in the output log format EVTX.
```

cluster1::> vserver audit modify -vserver vs1 -format evtx -events file-ops

```
[[ID443a290da62082b7445a521908dbd0ef]]
= vserver audit prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore the Audit configuration to Earlier Release of Data ONTAP

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver audit prepare-to-downgrade` command restores the Audit
configurations for ONTAP based on the input parameter disable-feature-set.

== Parameters

``-disable-feature-set <downgrade version>`` - Data ONTAP Version
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the ONTAP version that introduced the new Audit
features and needs to be removed. The value can be one of the following:
+
  * 9.0.0 - Disables the Audit features introduced in the ONTAP release
9.0.0. The following events are removed from the `event` list:
+
    * File share event. The corresponding parameter value is ``_file-
share_`` .
    * Audit policy change event. The corresponding parameter value is
``_audit-policy-change_`` .
    * Local user account management event. The corresponding parameter
value is ``_user-account_`` .
    * Local security group management event. The corresponding parameter
value is ``_security-group_`` .
    * Authorization policy change event. The corresponding parameter value
is ``_authorization-policy-change_`` .





== Examples
```

cluster1::*> vserver audit prepare-to-downgrade -disable-feature-set 9.0.0

```
[[ID6f34e5d6e61c2b048d17ed3e63d66569]]
= vserver audit rotate-log
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rotate audit log

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver audit rotate-log` command rotates audit logs for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which audit logs are
to be rotated. The Vserver audit configuration must already exist.
Auditing must be enabled for the Vserver.

== Examples

The following example rotates audit logs for Vserver vs1.
```

cluster1::> vserver audit rotate-log -vserver vs1

```
[[ID61ebef952927ca033d854253089e2d41]]
= vserver audit show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver audit show` command displays audit configuration information
about Vservers. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information about all the Vservers:

* Vserver name
* Audit state
* Target directory



You can specify the `-fields` parameter to specify which audit
configuration information to display about Vservers. + You can specify
additional parameters to display only information that matches those
parameters. For instance, to display information about the log file
rotation size of a Vserver whose value matches 10 MB, run the command with
the `-rotate-size 10MB` parameter.

You can specify the `-instance` parameter to display audit configuration
information for all Vservers in list form.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-log-save-details  ]``::
You can specify the `-log-save-details` parameter to display the following
information about all the Vservers:
+
  * Vserver name
  * Rotation file size
  * Rotation schedules
  * Rotation limit



| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information about the
specified Vserver.
```

```
``[-state {true|false}]`` - Auditing State::
If you specify this parameter, the command displays information about the
Vservers that use the specified audit state value.

``[-destination <text>]`` - Log Destination Path::
If you specify this parameter, the command displays information about the
Vservers that use the specified destination path.

``[-events {file-ops|cifs-logon-logoff|cap-staging|file-share|audit-
policy-change|user-account|authorization-policy-change|security-
group|async-delete}]`` - Categories of Events to Audit::
If you specify this parameter, the command displays information about the
Vservers that use the specified category of events that are audited. Valid
values are ``_file-ops_`` , ``_cifs-logon-logoff_`` , ``_cap-staging_`` ,
``_file-share_`` , ``_audit-policy-change_`` , ``_user-account_`` ,
``_security-group_`` and ``_authorization-policy-change_`` . ``_audit-
policy-change_`` will appear only in diag mode.

``[-format {xml|evtx}]`` - Log Format::
If you specify this parameter, the command displays information about the
Vservers that use the specified log format.

``[-rotate-size {<size>|-}]`` - Log File Size Limit::
If you specify this parameter, the command displays information about the
Vservers that use the specified log file rotation size.

``[-rotate-schedule-month <cron_month>,...]`` - Log Rotation Schedule:
Month::
If you specify this parameter, the command displays information about the
Vservers that use the specified month of the time-based log rotation
scheme. Valid values are January, February, March, April, May, June, July,
August, September, October, November, and December.

``[-rotate-schedule-dayofweek <cron_dayofweek>,...]`` - Log Rotation
Schedule: Day of Week::
If you specify this parameter, the command displays information about the
Vservers that use the specified day of the week of the time-based log
rotation scheme. Valid values are Sunday, Monday, Tuesday, Wednesday,
Thursday, Friday, and Saturday.

``[-rotate-schedule-day <cron_dayofmonth>,...]`` - Log Rotation Schedule:
Day::
If you specify this parameter, the command displays information about the
Vservers that use the specified day of the month of the time-based log
rotation scheme. Valid values range from 1 to 31.
```

```
``[-rotate-schedule-hour <cron_hour>,...]`` - Log Rotation Schedule:
Hour::
If you specify this parameter, the command displays information about the
Vservers that use the specified hour of the time-based log rotation
scheme. Valid values range from 0 (midnight) to 23 (11:00 p.m.).

``[-rotate-schedule-minute <cron_minute>,...]`` - Log Rotation Schedule:
Minute::
If you specify this parameter, the command displays information about the
Vservers that use the specified minute of the time-based log rotation
scheme. Valid values range from 0 to 59.

``[-rotate-schedule-description <text>]`` - Rotation Schedules::
If you specify this parameter, the command displays information about the
Vservers that use the specified rotation schedules. This field is derived
from the rotate-time fields.

``[-rotate-limit <integer>]`` - Log Files Rotation Limit::
If you specify this parameter, the command displays information about the
Vservers that use the specified rotation limit value.

``[-retention-duration
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Log Retention
Duration::
If you specify this parameter, the command displays information about the
Vservers audit logs retention duration.

``[-audit-guarantee {true|false}]`` - Strict Guarantee of Auditing::
If you specify this parameter, the command displays information about the
Vservers that have the specified audit guarantee value.

``[-charge-qos {true|false}]`` - Audit Log Extra Charge::
If you specify this parameter, the command displays information about the
Vservers that have the specified charge-qos value.

== Examples

The following example displays the name, audit state, event types, log
format, and target directory for all Vservers.
```

cluster1::> vserver audit show
Vserver State Event Types Log Format Target Directory
---------- ------ ---------- --------- ----------
vs1 false file-ops evtx /audit_log

```
The following example displays the Vserver names and details about the
audit log for all Vservers.
```

cluster1::> vserver audit show -log-save-details
Rotation Rotation
Vserver File Size Rotation Schedule Limit
---------- --------- ---------- --------
vs1 100MB - 0

```
The following example displays in list form all audit configuration
information about all Vservers.
```

cluster1::> vserver audit show -instance
Vserver: vs1
Auditing state: true
Log Destination Path: /audit_log
Categories of Events to Audit: file-ops
Log Format: evtx
Log File Size Limit: 100MB
Log Rotation Schedule: Month: -
Log Rotation Schedule: Day of Week: -
Log Rotation Schedule: Day: -
Log Rotation Schedule: Hour: -
Log Rotation Schedule: Minute: -
Rotation Schedules: -
Log Files Rotation Limit: 0
Log Retention Time: 0s

```
[[IDf3f0e54709e5f3de0d45e124ceebfc1b]]
= vserver audit audit-log-redirect create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create Audit Log Redirect Vserver destination

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver audit audit-log-redirect create` command creates an audit log
redirect configuration to designate a dedicated Vserver for all audit logs
in a cluster to reside in.

== Parameters

``-vserver <vserver name>`` - Audit Log Redirect Vserver [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify a dedicated Vserver for all audit logs to
reside in.

== Examples

The following example causes all audit logs in a cluster to be redirected
to ``_vserver1_`` :
```

cluster1::> vserver audit audit-log-redirect create -vserver vserver1

```
[[ID1fe70731eaf64d0e860857a062595f3d]]
= vserver audit audit-log-redirect delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete Audit Log Redirect Vserver destination

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver audit audit-log-redirect delete` command deletes the
dedicated Vserver for all audit logs in a cluster to reside in.

== Examples

The following example delete the audit-log-redirect Vserver.
```

cluster1::> vserver audit audit-log-redirect delete

```
[[IDbf9c7d87086b7101f534aee196a3b779]]
= vserver audit audit-log-redirect modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify Audit Log Redirect Vserver destination

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver audit audit-log-redirect modify` command changes the
dedicated Vserver for all audit logs in a cluster to reside in.

== Parameters

``[-vserver <vserver name>]`` - Audit Log Redirect Vserver [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify a dedicated Vserver for all audit logs to
reside in.

== Examples

The following example modify the current audit-log-redirect Vserver to
``_vserver1_`` :
```

cluster1::> vserver audit audit-log-redirect modify -vserver vserver1

```
[[ID15b5425591be3f41b4f5a8b0cc01d234]]
= vserver audit audit-log-redirect show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Audit Log Redirect Vserver destination

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver audit audit-log-redirect show` command displays the dedicated
Vserver for all audit logs in a cluster to reside in.

== Examples

The following example displays the audit-log-redirect Vserver:
```

cluster1::> vserver audit audit-log-redirect show
Audit Log Redirect Vserver: vserver1

```
:leveloffset: -1


= vserver check commands

:leveloffset: +1


[[IDe30fc198950329234e8163bb4a60ab3d]]
= vserver check lif-multitenancy run
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Run check for LIF multitenancy

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.
```

```
== Description


NOTE: This command is deprecated and may be removed in a future release of
ONTAP.



The run command checks the specified Vserver to verify that it has
connectivity to the configured external servers providing services such as
Active Directory, NIS, and DNS. The output can consist of three types of
messages. Failure messages indicate that a Vserver does not have the
connectivity required to a server exporting a service. Warning messages
indicate configuration or operational issues that are possible causes of
the failures. A success message is displayed if the Vserver has network
connectivity to each of the configured servers for each service.

You can use this command to verify configuration changes such as creating
a Vserver or changing the configured servers for one or more services. It
is also useful for diagnosing operational problems that result from
failures that could be caused by the inability to make network connections
to configured servers.

The services that are checked are DNS, NIS, CIFS preferred domain
controllers, CIFS discovered domain controllers, KDC, Active Directory,
Admin, Password, LDAP, and LDAP preferred Active Directory.

Only a single run for a Vserver is allowed to run in a cluster. If
multiple runs are attempted for a Vserver, a message will be displayed
indicating that a run is already in progress.

For each service, this command will ping each configured server until a
successful ping is completed. In certain circumstances where a subnet is
offline or LIFs are operationally down, this command may take a long time
to run. In order to show that forward progress is being made, an activity
indicator of a '.' is displayed for each ping sent.

The following fields are reported in table format. Some fields may not be
relevant to a type of message and will consist of the text "-".

* Vserver name
* Service external server is exporting
* Address of external server
* Connectivity to that external server
* More information describing the problem
* Suggestions to remediate the problems
```

```
* Success when there are no problems



== Parameters

``-vserver <vserver>`` - Vserver::
Use this parameter to specify the Vserver to check.

``[-verbose {true|false}]`` - Show Positive and Negative Result [.advaced-
priv]#(privilege: advanced)#::
When this parameter is specified the results of all connectivity tests
will be displayed in the success and failure cases.

== Examples

 This is an example of a successful run:
```

cluster1::> vserver check lif-multitenancy run -vserver vs0
..
SUCCESS: All external servers are reachable.

```
This is an example of a run with warnings and failures that need to be
corrected:
```

cluster1::> vserver check lif-multitenancy run -vserver vs0
Vserver Severity Service Address LIF Connected Details
--- --- --- --- --- --- ---
vs0 warning - - vs0_lif1 - operationally down
vs0 warning - - vs0_lif2 - operationally down
…
vs0 failure DNS 10.98.200.20 - no cache
…
vs0 failure NIS domain 10.98.13.53 - no cache
Error : command failed: FAILURES FOUND.
You must correct these failures to avoid service disruptions
in DOT 8.3 and above.
Corrective actions may include:
- removing decommissioned external servers from the vserver
configuration
- restoring network interfaces that are down
- adding network interfaces or routes
- modifying the locations where network interfaces may reside
(through
adjusting failover groups/policies or changing the home-node or
auto-revert settings).
For assistance, please consult the 8.3 Upgrade Document,
or contact support personnel.

```
At advanced privilege, additional information for messages at all
severities is displayed.
```

cluster1::*> vserver check lif-multitenancy run -vserver vs0 -verbose true
....
Vserver Severity Service Address LIF Connected Details
--- --- --- --- --- --- ---
vs0 info DNS 10.98.200.20 vs0_lif1 yes ping
....
vs0 info NIS domain 10.98.13.53 vs0_lif1 yes ping
SUCCESS: All external servers are reachable.

```
[[ID4a98762d5873de1842cea51064db9c9a]]
= vserver check lif-multitenancy show-results
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Show the results of the latest multitenancy network run

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future release of
ONTAP.

You can view detailed information about the latest completed run, or the
run for a Vserver.

* Vserver - name of vserver run was for
* Severity - severity of the message which is failure, warning, or info.
*
                  Failures are problems that need fixed.  Warnings are
potential
                  problems that may need to be fixed. Values are
"failure", "warning" or "info".

* Service - name of service that is being checked for connectivity
* Address - address of server configured for the above service that is
being
*
```

```
                  checked for connectivity.

* LIF - the LIF a successful connectivity check to the above server was
made from
* Connected - true of there is connectivity, false if there is not
* Status - additional information useful for resolving issues



== 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.

``[-vserver <vserver>]`` - Vserver::
Selects the messages matching the specified Vserver

``[-severity <text>]`` - Severity::
Selects the messages matching the specified severity of failure, warning,
and info.

``[-service <text>]`` - Service Name::
Selects the messages matching the specified service.

``[-address <text>]`` - Address of Server::
Selects the messages matching the specified address.

``[-lif <lif-name>]`` - Logical Interface::
Selects the messages matching the specified LIF.

``[-connected {yes|no}]`` - Vserver Connectivity::
Selects the messages matching the specified connectivity.

``[-status <text>]`` - Additional Information::
Selects the messages matching the specified search criteria.

== Examples

Runs that are successful will not have any content.
```

cluster1::> vserver check lif-multitenancy show-results -vserver vs0
This table is currently empty.

```
Successful runs made with -verbose true will show the LIF used to Ping the
nework address from.
```

cluster1::> vserver check lif-multitenancy show-results -vserver vs0
Network Logical
Vserver Severity Service Address Interface Connected Status
---------- -------- ----------- ----------- --------- ---------- -------
vs0
info DNS 10.98.200.20
vs0_lif1 yes ping
info NIS domain 10.98.13.53 vs0_lif1 yes ping
2 entries were displayed.

```
Runs that fail display each failure that needs to be fixed.
```

cluster1::> vserver check lif-multitenancy show-results -vserver vs0
Network Logical
Vserver Severity Service Address Interface Connected Status
---------- -------- ----------- ----------- --------- ---------- -------
vs0
warning - - vs0_lif1 - operationally down
warning - - vs0_lif2 - operationally down
failure DNS 10.98.200.20
- no cache
failure NIS domain 10.98.13.53 - no cache
4 entries were displayed.

```
[[IDfeba911b737249e945515b07871a869d]]
= vserver check lif-multitenancy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Show the summary of the latest multitenancy network run

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future release of
```

```
ONTAP.

You can view summary information about the latest completed run, or the
run in progress for a Vserver. It will show the following fields:

* Vserver - Name of Vserver that was checked for LIF connectivity
* Start Time - Date And Time the run was started
* Status - Not Started, In Progress, Complete, or Aborted
* Success - Yes if the run has a Status of Complete with no failures. No
if the run has a status of Complete with one or more failures.
* Updated - The date and time the scan was last updated.



== 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.

``[-vserver <vserver>]`` - Vserver::
Selects the summary information matching the specified Vserver.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Start Time::
Selects the summary information matching the specified date and time the
run was started

``[-status {not started|in progress|complete|aborted}]`` - Run Status::
Selects the summary information matching the specified status of the run.

``[-success {yes|no}]`` - Successful Run::
Selects the summary information matching the specified success or failure
of the run.

``[-updated <MM/DD/YYYY HH:MM:SS>]`` - Run Updated::
Selects the summary information matching the last time the run was still
in progress.

== Examples

This is what a successful run looks like:
```

cluster1::> vserver check lif-multitenancy show
Vserver Start Time Status Success
--- --- --- ---
vs0
7/16/2014 14:28:35 complete yes

```
This is what a failed run looks like:
```

cluster1::> vserver check lif-multitenancy show
Vserver Start Time Status Success
--- --- --- ---
vs0
7/16/2014 14:40:55 complete no

```
This is what specifying the Vserver looks like:
```

cluster1::> vserver check lif-multitenancy show -vserver vs0
Vserver: vs0
Start Time: 7/16/2014 14:40:55
Run Status: complete
Successful Run: no

```
Advanced privilege adds in the Updated field.
```

cluster1::*> vserver check lif-multitenancy show
Vserver Start Time Status Success Updated
--- --- --- --- ---
vs0
7/16/2014 14:40:55
complete no 7/16/2014 14:40:56

```
:leveloffset: -1


= vserver cifs commands

:leveloffset: +1


[[ID64e6d23e900e9251f67ab05ac0174244]]
= vserver cifs add-netbios-aliases
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add NetBIOS aliases for the CIFS server name

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 The `vserver cifs add-netbios-aliases` command creates or adds a list of
NetBIOS aliases for the CIFS server name.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which NetBIOS alias
are to be created or added.

``-netbios-aliases <NetBIOS>,...`` - List of NetBIOS Aliases::
This parameter specifies one or more NetBIOS aliases to be added to an
existing list of NetBIOS aliases. A new list of NetBIOS aliases is created
if the list is currently empty.

== Examples

The following example creates a new list of NetBIOS aliases for Vserver
vs_a.
```

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a
Server Name: CIFS_SERVER

NetBIOS Aliases: -

cluster1::> cifs add-netbios-aliases -netbios-aliases alias_1,alias_2,alias_3

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a
Server Name: CIFS_SERVER
NetBIOS Aliases: ALIAS_1, ALIAS_2, ALIAS_3

```
The following example adds several NetBIOS aliases for the CIFS server
CIFS_SERVER on Vserver vs_a.
```

cluster1::> cifs add-netbios-aliases -netbios-aliases alias_4,alias_5,alias_6

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a

```
    Server Name: CIFS_SERVER
NetBIOS Aliases: ALIAS_1, ALIAS_2, ALIAS_3, ALIAS_4,
                 ALIAS_5, ALIAS_6
```

cluster1::> vserver cifs add-netbios-aliases -vserver v1 -netbios-aliases alias_7

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a

```
    Server Name: CIFS_SERVER
NetBIOS Aliases: ALIAS_1, ALIAS_2, ALIAS_3, ALIAS_4,
                 ALIAS_5, ALIAS_6, ALIAS_7
```

```
[[ID38bdba9344f3e38173186d504e67eb5c]]
= vserver cifs check
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Validation Status of CIFS Configuration from Each Node

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

Use the `vserver cifs check` command to check the status of configured
CIFS server on a particular vserver.

== 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.

``[-vserver <vserver name>]`` - Vserver::
Use this parameter to specify the Vserver whose CIFS server needs to be
validated.

``[-node {<nodename>|local}]`` - Node::
Use this parameter to specify the node name from which CIFS server
connectivity needs to be validated.

``[-netbios-name <TextNoCase>]`` - CIFS NetBIOS Name::
Use this parameter to display netbios-name of the configured CIFS server.

``[-cifs-status <TextNoCase>]`` - CIFS Server Status::
Use this parameter to display status of configured CIFS server.

``[-site <TextNoCase>]`` - CIFS Server Site::
This parameter specifies the site discovered from ONTAP for the Active
Directory domain associated with the CIFS server. If the discovery fails,
this parameter will be updated with the default-site of associated cifs
server.

``[-server <TextNoCase>]`` - Domain Controller Name::
Use this parameter to display Domain name of the configured CIFS server.

``[-server-ip <text>]`` - Domain Controller IP Addr::
Use this parameter to display IP-address of the configured CIFS server.

``[-status {down|up}]`` - Connectivity Status::
Use this parameter to display information only about CIFS servers with a
status that matches the value you specify.

``[-status-details <text>]`` - Connectivity Status Details::
```

```
Use this parameter to display information only about CIFS servers with
status details that match the value you specify.

== Examples

The following example checks the connectivity of CIFS server on vserver
vs0 from each node.
```

cluster1::> vserver cifs check -vserver vs0
Vserver : vs0
Cifs NetBIOS Name : NEWSERVER
Cifs Status : up
Site : Bangalore

Node Name DC Server Name DC Server IP Status Status Details
--------- ---------- ---------- ------- ----------
node1 CIFSSERVER.COM 10.11.12.13 up Response time (msec): 55
node2 CIFSSERVER.COM 10.11.12.13 up Response time (msec): 70
node3 CIFSSERVER.COM 10.11.12.13 down Secd: No Server available.

```
[[IDbaeeb7495caab283418cd853864cd907]]
= vserver cifs create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a CIFS server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs create` command creates a CIFS server on a Vserver. When
you create the CIFS server, you can add it to an existing CIFS domain, or
you can join it to a workgroup. When you add it to an existing CIFS
domain, the storage system prompts you to provide the credentials of a
user account that has sufficient privileges to add computers to the `-ou`
container within the `-domain` domain. The user account must have a
password that cannot be empty. If the new CIFS server is joining a domain,
this command might take several minutes to complete.

NOTE: Each Vserver can have only one CIFS server.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which to create the
CIFS server. The Vserver must already exist.

``-cifs-server <NetBIOS>`` - CIFS Server NetBIOS Name::
This parameter specifies the name of the CIFS server (up to 15
characters).

{ ``-domain <TextNoCase>`` - Fully Qualified Domain Name::
This parameter specifies the name of the Active Directory domain to
associate with the CIFS server.

``[-ou <text>]`` - Organizational Unit::
This parameter specifies the organizational unit within the Active
Directory domain to associate with the CIFS server. By default, this
parameter is set to CN=Computers. When specifying this parameter, specify
only the organizational unit portion of the distinguished name. ONTAP
appends the value provided for the required `-domain` parameter onto the
value provided for `-ou` parameter to produce the Active Directory
distinguished name, which is used to associate with the CIFS server.
+
NOTE: Nested OUs must be provided in a specific order with all containers
separated by a comma. Reading from left to right you travel up the
directory tree until you reach the root OU.
+


``[-default-site <text>]`` - Default Site Used by LIFs Without Site
Membership::
This parameter specifies the site within the Active Directory domain to
associate with the CIFS server if ONTAP cannot determine an appropriate
site. This parameter will also be used for discovering KDCs in the trusted
domain if ONTAP cannot determine an appropriate site in the trusted
domain.

| ``-workgroup <NetBIOS>`` - Workgroup Name }::
This parameter specifies the name of the workgroup (up to 15 characters).

``[-keytab-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}]``
- Kerberos Keytab File URI [.advaced-priv]#(privilege: advanced)#::
This parameter specifies loading a keytab file from the specified URI.
This is applicable if the CIFS server is being created in realm mode or
domain mode.
```

```
``[-status-admin {down|up}]`` - CIFS Server Administrative Status::
Use this parameter to specify whether the initial administrative status of
the cifs server is up or down. The default setting is `up` .

``[-comment <text>]`` - CIFS Server Description::
This optional parameter specifies a text comment for the server. CIFS
clients can see this CIFS server description when browsing servers on the
network. The comment can be up to 256 characters long. If there is a space
in the descriptive remark or the path, you must enclose the entire string
in quotation marks.

``[-netbios-aliases <NetBIOS>,...]`` - List of NetBIOS Aliases::
This parameter specifies a list of NetBIOS aliases, which are alternate
names to the CIFS server name.

== Examples

The following example creates a CIFS server CIFSSERVER1 for Vserver vs1
and domain EXAMPLE.com.
```

cluster1::> vserver cifs create -vserver vs1 -cifs-server CIFSSERVER1 -domain EXAMPLE.com
In order to create an Active Directory machine account for the CIFS server, you
must supply the name and password of a Windows account with sufficient
privileges to add computers to the "CN=Computers" container within the
"EXAMPLE.com" domain.

Enter the user name: Administrator

Enter the password:

```

```

cluster1::> vserver cifs create -vserver vs1 -cifs-server CIFSSERVER2 -domain EXAMPLE.com -keytab-uri http://nbsweb.eng.btc.netapp.in/~user/Sample1.keytab
Info: SMB1 protocol version is disabled on this CIFS server. If required, use
the (privilege: advanced) command "vserver cifs options modify -vserver vs1
-smb1-enabled true" to enable it.

```
The following example creates a CIFS server CIFSSERVER1 for Vserver vs1
and workgroup Sales:
```

cluster1::> vserver cifs create -vserver vs1 -cifs-server CIFSSERVER1 -workgroup Sales

```
The following example creates a CIFS server CIFSSERVER1 for Vserver vs1
and domain EXAMPLE.com with a user Administrator1 from a different domain,
in this case an administrator from a trusted domain TRUST.LAB.COM:
```

cluster1::> vserver cifs create -vserver vs1 -cifs-server CIFSSERVER1 -domain EXAMPLE.com
In order to create an Active Directory machine account for the CIFS server, you
must supply the name and password of a Windows account with sufficient
privileges to add computers to the "CN=Computers" container within the
"EXAMPLE.com" domain.

Enter the user name: Administrator1@TRUST.LAB.COM

Enter the password:

```
The following example creates a CIFS server CIFSSERVER1 for Vserver vs1
with domain EXAMPLE.com using nested OUs:
```

cluster1::> vserver cifs create -vserver vs1 -cifs-server CIFSSERVER1 -domain EXAMPLE.com -ou
OU=developers,OU=engineering,OU=corp
In order to create an Active Directory machine account for the CIFS server, you
must supply the name and password of a Windows account with sufficient
privileges to add computers to the "OU=developers,OU=engineering,OU=corp" container within the
"EXAMPLE.com" domain.

Enter the user name: Administrator

Enter the password:

```
[[IDf608c7f50c4bf1b0a42a6cccaa358c47]]
= vserver cifs delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a CIFS server.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs delete` command deletes a CIFS server.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver for the CIFS server you want to
delete.

``[-skip-ad-account-delete {true|false}]`` - If this option is enabled,
the local CIFS configuration will be removed without deleting the AD
account. The default setting for this field is false.::
This parameter skips the deletion of the AD account and proceeds directly
to the deletion of the local CIFS configuration.

``[-client-id <text>]`` - Application (Client) ID of Deployed Azure
Application::
This parameter specifies the Application (Client) ID of Deployed "Entra
ID" Application.

``[-tenant-id <text>]`` - Directory (Tenant) ID of Deployed Azure
Application::
This parameter specifies the Directory (Tenant) ID of Deployed "Entra ID"
Application.

``[-client-certificate <text>]`` - Certificate for Azure Application::
This parameter specifies the Certificate to authenticate with Deployed
"Entra ID" Application.

== Examples

The following example deletes the CIFS server from a Vserver named vs1:
```

cluster1::> vserver cifs delete -vserver vs1

```
[[IDc0ef9a5ec681040992df5043816a6550]]
= vserver cifs modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a CIFS server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs modify` command modifies the site within the Active
Directory domain to associate with the CIFS server if ONTAP cannot
determine an appropriate site. You also can modify the name and ou of the
CIFS server, join to a new domain or a workgroup, or rejoin to current
domain. When a CIFS server is joining a domain, this command might take
several minutes to complete.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver for the CIFS server whose associated
site you want to modify.

``[-cifs-server <NetBIOS>]`` - CIFS Server NetBIOS Name::
This parameter specifies the name of the CIFS server (up to 15
characters). Before setting this parameter, the CIFS server must be
stopped using the `vserver cifs modify``-status-admin```_down_`` command.
When the command completes successfully, the administrative status of the
CIFS server is automatically set to ``_up_`` .

{ ``[-domain <TextNoCase>]`` - Fully Qualified Domain Name::
This parameter specifies the fully qualified name of the Active Directory
domain to associate with the CIFS server. Before setting this parameter,
the CIFS server must be stopped using the `vserver cifs modify``-status-
admin```_down_`` command. When the command completes successfully, the
administrative status of the CIFS server is automatically set to ``_up_``
.

``[-ou <text>]`` - Organizational Unit::
```

```
This parameter specifies the organization unit within the Active Directory
domain to associate with the CIFS server. By default, this parameter is
set to CN=Computers. Before setting this parameter, the CIFS server must
be stopped using the `vserver cifs modify``-status-admin```_down_``
command. When the command completes successfully, the administrative
status of the CIFS server is automatically set to ``_up_`` . Modifications
to this parameter are not supported for workgroup CIFS servers.

``[-default-site <text>]`` - Default Site Used by LIFs Without Site
Membership::
This parameter specifies the site within the Active Directory domain to
associate with the CIFS server if ONTAP cannot determine an appropriate
site. Modifications to this parameter are not supported for workgroup CIFS
servers.

| ``[-workgroup <NetBIOS>]`` - Workgroup Name }::
This parameter specifies the name of the workgroup (up to 15 characters).

``[-keytab-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}]``
- Kerberos Keytab File URI [.advaced-priv]#(privilege: advanced)#::
This parameter specifies loading a keytab file from the specified URI.
This is applicable if the CIFS server is being created in realm mode or
domain mode.

``[-status-admin {down|up}]`` - CIFS Server Administrative Status::
Use this parameter to modify the administrative status of the cifs server.
Modify the administrator status to `down` to stop cifs access.

``[-comment <text>]`` - CIFS Server Description::
Use this parameter to modify the comment of the server.

== Examples

The following example changes the default site and administrative status
of the CIFS server associated with Vserver "vs1":
```

cluster1::> vserver cifs modify -vserver vs1 -default-site default -status-admin up

```
The following example modifies the Active Directory domain and ou for the
CIFS server associated with Vserver "vs1". The administrative status of
the CIFS server must be set to "down" to proceed with Active Directory
domain modification. If the command completes successfully, the
administrative status is automatically set to "up".
```

cluster1::> vserver cifs modify -vserver vs1 -domain example.com -ou ou=example_ou -cifs-server example

-status-admin down

In order to create an Active Directory machine account for the CIFS server, you
must supply the name and password of a Windows account with sufficient
privileges to add computers to the "ou=example_ou" container within the "example.com"
domain.

Enter the user name: administrator

Enter the password:

cluster1::>

```
The following example modifies the CIFS server associated with Vserver
"vs1" from a domain to a workgroup. The administrative status of the CIFS
server must be set to "down" for this command. If the command completes
successfully, the administrative status is automatically set to "up".
```

cluster1::> vserver cifs modify -vserver vs1 -workgroup Sales -status-admin down

Warning: To enter workgroup mode, all domain-based features must be disabled
and their configuration removed automatically by the system,
including continuously-available shares, shadow copies, and AES.
However, domain-configured share ACLs such as
"EXAMPLE.COM\userName" will not work properly, but cannot be
removed by ONTAP. Remove these share ACLs as soon as possible
using external tools after the command completes. If AES is enabled,
you may be asked to supply the name and password of a Windows account
with sufficient privileges to disable it in the "EXAMPLE.COM" domain.
Do you want to continue? {y|n}: y

cluster1::>

```
The following example modifies the CIFS server associated with Vserver
"vs1" from a workgroup to a domain. The administrative status of the CIFS
server must be set to "down" for this command. If the command completes
successfully, the administrative status is automatically set to "up".
```

cluster1::> vserver cifs modify -vserver vs1 -domain example.com -status-admin down

In order to create an Active Directory machine account for the CIFS server, you
must supply the name and password of a Windows account with sufficient
privileges to add computers to the "ou=example_ou" container within the "example.com"
domain.

Enter the user name: administrator

Enter the password:

cluster1::>

```
The following example modifies the CIFS server associated with Vserver
"vs1" from a workgroup to a domain using keytab-uri. The administrative
status of the CIFS server must be set to "down" for this command. If the
command completes successfully, the administrative status is automatically
set to "up".
```

cluster1::> vserver cifs modify -vserver vs1 -domain example.com -keytab-uri http://nbsweb.eng.btc.netapp.in/~shravanp/Sample1.keytab -status-admin down

cluster1::>

```
The following example modifies the CIFS server name associated with
Vserver "vs1" from above example. The administrative status of the CIFS
server must be set to "down" to proceed with Active Directory domain
modification. If the command completes successfully, the administrative
status is automatically set to "up" and there will be a job running to
update related configurations.
```

cluster1::> vserver cifs modify -vserver vs1 -cifs-server new_example -status-admin down

In order to create an Active Directory machine account for the CIFS server, you
must supply the name and password of a Windows account with sufficient
privileges to add computers to the "ou=example_ou" container within the "example.com"
domain.

Enter the user name: administrator

Enter the password:

Successfully queued CIFS Server Modify job [id: xx] for CIFS server
"NEW_EXAMPLE". To view the status of the job, use the "job show -id <jobid>"
command.

cluster1::>

```
[[ID809755e24e1f16a33ddd694270da0df9]]
= vserver cifs nbtstat
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NetBIOS information over TCP connection

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The vserver cifs nbtstat command displays information about NetBIOS over
TCP (NBT) connections for the cluster. It displays the IP address
associated with the interfaces, the IP addresses of the WINS servers in
use, and information about the registered NetBIOS names for the cluster.
You can use this command to troubleshoot NetBIOS name resolution problems.

NOTE: NetBIOS name service (NBNS) over IPv6 is not supported.



== 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>|local}]`` - Node::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified Vserver.

``[-nbt-name <text>]`` - NBT Name::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified NetBIOS name.

``[-netbios-suffix <Hex String>]`` - NetBIOS Suffix::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified NetBIOS suffix.

``[-interface <IP Address>,...]`` - Interfaces::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified IP address.

``[-wins-servers <IP Address>,...]`` - Servers::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified WINS servers.
```

```
``[-server-state <text>,...]`` - Server State (active, inactive)::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified WINS server state. The
following are possible values for this parameter:
+
  * active
  * inactive



``[-nbt-scope <text>]`` - NBT Scope::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified NetBIOS name scope.

``[-nbt-mode <text>]`` - NBT Mode::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified NetBIOS name service mode.
The following are possible values for this parameter:
+
  * 'p' - Point to Point
  * 'h' - Hybrid
  * 'm' - Mixed
  * 'b' - Broadcast



``[-state <text>]`` - State::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified NetBIOS name registration
state. The following are possible values for this parameter:
+
  * must_register
  * must_unregister
  * wins
  * broadcast
  * name_released
  * wins_conflict
  * broadcast_conflict



``[-time-left <integer>]`` - Time Left::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified registration time left in
minutes with the WINS server.
```

```
``[-type <text>]`` - Type::
If you specify this optional parameter, the command displays the NetBIOS
name service information only for the specified name registration type.
The following are possible values for this parameter:
+
  * registered
  * active
  * permanent
  * group



== Examples

The following example displays the NetBIOS name service information.
```

cluster1::> nbtstat
(vserver cifs nbtstat)

```
          Vserver: vs1
          Node:    cluster1-01
          Interfaces:
                   10.10.10.32
                   10.10.10.33
         Servers:
                   17.17.1.2  (active  )
         NBT Scope:
                  [ ]
         NBT Mode:
                  [h]
         NBT Name           NetBIOS Suffix   State         Time Left
Type
         ------------------ ---------------  ------------- ---------
-----
         CLUSTER_1          00               wins          57
         CLUSTER_1          20               wins          57
Vserver: vs1
         Node:    cluster1-02
         Interfaces:
                 10.10.10.35
         Servers:
                 17.17.1.2  (active  )
         CLUSTER_1          00               wins          58
         CLUSTER_1          20               wins          58
         4 entries were displayed.
```

```
[[ID886e9c4cffc7caaf79f24eb39ce9884c]]
= vserver cifs prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore the CIFS Configurations to Earlier Release of Data ONTAP Version

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver cifs prepare-to-downgrade` command restores the CIFS
configurations for ONTAP based on the input parameter disable-feature-set.

== Parameters

``-disable-feature-set <downgrade version>`` - Data ONTAP Version
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the ONTAP release for which the CIFS
configurations are restored. The value can be one of the following:
+
  * 8.3.1 - Restores the CIFS configurations for ONTAP release 8.3.1.
These features include:
+
    * FPolicy "close with read" filters from FPolicy events.
    * CIFS server options `-guest-unix-user` and `-is-admin-users-mapped-
to-root-enabled` .
    * CIFS security option `is-smb-encryption-required` .
    * Storage-Level Access Guard (SLAG) for qtrees.
    * CIFS share property `encrypt-data.`


  * 8.3.2 - Restores the CIFS configurations for ONTAP release 8.3.2.
These features include:
+
    * CIFS server option `-grant-unix-group-perms-to-others` .


  * 9.0.0 - Restores the CIFS configurations for ONTAP release 9.0.0.
These features include:
+
    * Disable CIFS multichannel feature and close all multichannel
```

```
connections.
    * Delete all the name-mapping entries that have a hostname or an
address field configured.
    * Terminate all SMB 3.1 client connections.
    * Terminate all client connections that have large MTU negotiated.
    * Remove the symlink property `no-strict-security` .
    * Remove all symlink pathmap entries with locality `freelink` .





== Examples
```

cluster1::*> vserver cifs prepare-to-downgrade -disable-feature-set 8.3.1

```
```

cluster1::*> vserver cifs prepare-to-downgrade -disable-feature-set 8.3.2

```
```

cluster1::*> vserver cifs prepare-to-downgrade -disable-feature-set 9.0.0

```
[[IDbb23909df0724ce45138d5eba27ec0e7]]
= vserver cifs remove-netbios-aliases
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove NetBIOS aliases

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 The `vserver cifs remove-netbios-aliases` command deletes NetBIOS aliases
for the CIFS server.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver from which the list of
NetBIOS aliases are deleted.

``-netbios-aliases <NetBIOS>,...`` - List of NetBIOS Aliases::
This parameter specifies one or more NetBIOS aliases to be deleted. To
delete all the NetBIOS aliases of a Vserver use '-'.

== Examples

The following example deletes NetBIOS aliases for the CIFS server
CIFS_SERVER on Vserver vs_a.
```

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a
Server Name: CIFS_SERVER
NetBIOS Aliases: ALIAS_1, ALIAS_2, ALIAS_3, ALIAS_4,
ALIAS_5, ALIAS_6, ALIAS_7

cluster1::> cifs remove-netbios-aliases -netbios-aliases alias_1,alias_3,alias_5

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a
Server Name: CIFS_SERVER
NetBIOS Aliases: ALIAS_2, ALIAS_4, ALIAS_6, ALIAS_7

cluster1::> cifs remove-netbios-aliases -netbios-aliases alias_7

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a
Server Name: CIFS_SERVER
NetBIOS Aliases: ALIAS_2, ALIAS_4, ALIAS_6

cluster1::> cifs remove-netbios-aliases -netbios-aliases -

cluster1::> cifs show -display-netbios-aliases

Vserver: vs_a
Server Name: CIFS_SERVER
NetBIOS Aliases: -

```
[[ID4539453233019b1d876778b5dddec3db]]
= vserver cifs repair-modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Repair a partially-failed Vserver CIFS server modify operation

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

Use this `vserver cifs repair-modify -vserver <vserver name>` command when
the background job created during a Vserver CIFS server modify operation
fails.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies a Vserver containing a configured CIFS server
that has been modified.

== Examples

The following example starts the CIFS server modify job on Vserver vs1
successfully:
```

cluster1::*> vserver cifs repair-modify -vserver vs1

Successfully queued CIFS Server Modify job [id: 10] for CIFS server "CIFSNAME1".
To view the status of the job, use the "job show -id <jobid>" command.

cluster1::*>

```
The following example fails the command with specific error:
```

cluster1::*> vserver cifs repair-modify -vserver vs2

Error: Job Out of memory. Failed to queue CIFS Server Modify Job for CIFS
server "CIFSNAME2". Retry the operation by running (privilege: advanced)
"vserver cifs repair-modify -vserver vs2".
Error: command failed: unable to save data

cluster1::*>

```
[[IDbd9d96daff4a17e0f9ad369952411672]]
= vserver cifs security-encryption-required-dc-connections-prepare-to-
downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disabled encryption-required-for-dc-connections option and capability for
downgrade.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver cifs security-encryption-required-dc-connections-prepare-to-
downgrade` prepares the cluster for downgrade by disabling
SMB3.encryption.required.dc.connection capability

== Examples
```

cluster1::*> vserver cifs security-encryption-required-dc-connections-prepare-to-downgrade

```
[[IDeeeebb898a9577688f6852a9c7be5ad0]]
= vserver cifs show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display CIFS servers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs show` command displays information about CIFS servers.
The command output depends on the parameter or parameters specified with
the command. If you do not specify any parameters, the command displays
the following information about all CIFS servers:

* Vserver name
* CIFS server NetBIOS name
* Domain or workgroup name
* Authentication style



You can specify the `-fields` parameter to specify which fields of
information to display about CIFS servers. In addition to the fields
above, you can display the following fields:

* Default site
* Fully-qualified domain name

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
CIFS servers that are in the CIFS domain named RUBY, run the command with
the `-domain-workgroup RUBY` parameter.

You can specify the ` -instance` parameter to display all information for
all CIFS servers in list form.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-display-netbios-aliases  ]``::
If you specify this parameter, the command displays information about
configured NetBIOS aliases.
```

```
| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the CIFS servers for the specified Vserver.

``[-cifs-server <NetBIOS>]`` - CIFS Server NetBIOS Name::
If you specify this parameter, the command displays information only for
CIFS servers that match the specified CIFS server NetBIOS name.

``[-domain-workgroup <CIFS domain>]`` - NetBIOS Domain/Workgroup Name::
If you specify this parameter, the command displays information only for
CIFS servers that are in the specified NetBIOS domain or workgroup.

``[-domain <TextNoCase>]`` - Fully Qualified Domain Name::
If you specify this parameter, the command displays information only for
CIFS servers that are in the specified domain.

``[-ou <text>]`` - Organizational Unit::
If you specify this parameter, the command displays information only for
CIFS servers that are in the specified organizational unit.

``[-default-site <text>]`` - Default Site Used by LIFs Without Site
Membership::
If you specify this parameter, the command displays information only for
CIFS servers that have the specified default site.

``[-workgroup <NetBIOS>]`` - Workgroup Name::
If you specify this parameter, the command displays information only for
CIFS servers that are in the specified workgroup.

``[-auth-style {domain|workgroup|realm}]`` - Authentication Style::
If you specify this parameter, the command displays information only for
CIFS servers that match the specified authentication style.

``[-status-admin {down|up}]`` - CIFS Server Administrative Status::
If you specify this parameter, the command displays information only for
CIFS servers that match the specified administrative status.

``[-comment <text>]`` - CIFS Server Description::
If you specify this parameter, the command displays information only for
CIFS servers that match the specified comment field.

``[-netbios-aliases <NetBIOS>,...]`` - List of NetBIOS Aliases::
If you specify this parameter, the command displays information only for
```

```
CIFS servers that have specified NetBIOS alias.

== Examples

The following example displays a subset of the information about all CIFS
servers:
```

cluster1::> vserver cifs show
Server Domain/Workgroup
Vserver Name Name Authentication Style
---------- ---------- ---------- ----------
vs1 CIFSSERVER1 EXAMPLE domain

```
The following example displays all information about all CIFS-enabled
Vservers in list form:
```

cluster1::> vserver cifs show -instance
Vserver: vs1
CIFS Server NetBIOS Name: CIFSSERVER1
NetBIOS Domain/Workgroup Name: EXAMPLE
Fully Qualified Domain Name: EXAMPLE.COM
Organizational Unit: CN=Computers
Default Site Used by LIFs Without Site Membership:
Workgroup Name: -
Authentication Style: domain
CIFS Server Administrative Status: up
CIFS Server Description:
List of NetBIOS Aliases: ALIAS_2, ALIAS_4, ALIAS_6

```
The following example displays the NetBIOS aliases for the CIFS server
CIFSSERVER1
```

cluster1::> cifs show -display-netbios-aliases

Vserver: vs1
Server Name: CIFSSERVER1
NetBIOS Aliases: ALIAS_2, ALIAS_4, ALIAS_6

```
[[IDd7a450f7573dcc55640eeee7d9ff452d]]
= vserver cifs start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a CIFS server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command starts the CIFS server on the specified Vserver. The CIFS
server must already exist. To create a CIFS server, run
xref:{relative_path}vserver-cifs-create.html[vserver cifs create] .

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies a Vserver containing a configured CIFS server
that has been stopped.

== Examples

The following example starts the CIFS server on Vserver vs1:
```

cluster1::> cifs start -vserver vs1

```
== Related Links

* xref:{relative_path}vserver-cifs-create.html[vserver cifs create]




[[IDa5ac7f63088baa6bc06cc093234102c0]]
= vserver cifs stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop a CIFS server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command stops the CIFS server on the specified Vserver.

NOTE: Established sessions will be terminated and their open files closed.
Workstations with cached data will not be able to save those changes,
which could result in data loss.



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies a Vserver containing a configured CIFS server
that is running.

== Examples

The following example stops the CIFS server on Vserver vs1:
```

cluster1::> cifs stop -vserver vs1

```
[[IDe505098306b3b92f1d75b6fe4790f197]]
= vserver cifs branchcache create
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create the CIFS BranchCache service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs branchcache create` command creates the configuration
for computing and retrieving BranchCache hash data. Only a single instance
of the BranchCache service can be created on a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver on which you want to set
up the BranchCache service.

``[-versions {v1-enable|v2-enable|enable-all}]`` - Supported BranchCache
Versions::
This optional parameter specifies a list of versions of the BranchCache
protocol that the storage system supports. The default is `enable-all` .
This list can include one or more of the following:
+
  * v1-enable - This option enables BranchCache Version 1.
  * v2-enable - This option enables BranchCache Version 2.
  * enable-all - This option enables all supported versions of
BranchCache.



``-hash-store-path <text>`` - Path to Hash Store::
This parameter specifies an existing directory into which the hash data is
stored. Read-only paths, such as snapshot directories, are not allowed.

``[-hash-store-max-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Size of
the Hash Store::
This optional parameter specifies the maximum size to use for the hash
data. If the size of the hash data exceeds this value, older hashes are
deleted to make room for newer hashes. The default is 1 GB.

``[-server-key <text>]`` - Encryption Key Used to Secure the Hashes::
This optional parameter specifies a server key that the BranchCache
```

```
service uses to prevent clients from impersonating the BranchCache server.

``[-operating-mode <BranchCache Mode>]`` - CIFS BranchCache Operating
Modes::
This optional parameter specifies the mode in which the BranchCache
service operates. The default is `per-share` . Possible values include:
+
  * disable - This option disables the BranchCache service for the
Vserver.
  * all-shares - This option enables the BranchCache service for all the
shares on this Vserver.
  * per-share - This option enables the BranchCache service on a per-share
basis. You can enable the BranchCache service on an existing share by
adding the ``_branchcache_`` flag in the `-share-properties` parameter of
the xref:{relative_path}vserver-cifs-share-modify.html[vserver cifs share
modify] command.



== Examples

The following example creates the BranchCache service on the Vserver named
vs1. The path to the hash store is /vs1_hash_store.
```

cluster1::> vserver cifs branchcache create -vserver vs1 -hash-store-path /vs1_hash_store

```
 The following example creates the BranchCache service on the Vserver vs1.
The path to the hash store is /vs_hash_store. The service is enabled on
all the shares of the Vserver, supports BranchCache version 2, supports a
maximum of 1 GB of BranchCache hashes, and secures the hashes using the
key "vs1 secret".
```

cluster1::> vserver cifs branchcache create -vserver vs1 -hash-store-path /vs1_hash_store -operating-mode all-shares -versions v2-enable -hash-store-max-size 1GB -server-key "vs1 secret"

```
== Related Links

* xref:{relative_path}vserver-cifs-share-modify.html[vserver cifs share
modify]




[[IDb9b1b87285cd711d4962eb2698275d7a]]
= vserver cifs branchcache delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop and remove the CIFS BranchCache service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs branchcache delete` command stops and removes the
Vserver BranchCache configuration.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver whose BranchCache
configuration you want to remove.

``-flush-hashes {true|false}`` - Delete Existing Hashes::
This parameter specifies whether to keep or delete all existing hashes
after deleting the BranchCache service.

== Examples

The following example stops and removes the BranchCache service on the
Vserver vs1. It also deletes all existing hashes.
```

cluster1::> vserver cifs branchcache delete -flush-hashes true -vserver vs1

```
[[IDcfe3d808c881731b26979f56c75a7b0e]]
= vserver cifs branchcache hash-create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Force CIFS BranchCache hash generation for the specified path or file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs branchcache hash-create` command causes the BranchCache
service to compute hashes for a single file, for a directory, or for all
the files in a directory structure if you specify the `-recurse` option.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver on which the hash is
computed.

``-path <text>`` - Path of File or Directory to Hash::
This parameter specifies the path of the directory or file for which
hashes are to be computed. If a file is specified, the hashes are computed
on the whole file. If a directory is specified, hashes are computed on all
files within the directory.

``-recurse {true|false}`` - Process All Files in the Directory
Recursively::
If this option is set to true and the `-path` parameter specifies a
directory, hashes are computed recursively for all directories in the
path.

== Examples

The following example creates hashes for the file "report.doc":
```

cluster1::> vserver cifs branchcache hash-create -vserver vs1 -path /repository/report.doc -recurse false

```
The following example creates hashes for all the files in the directory
"repository":
```

cluster1::> vserver cifs branchcache hash-create -vserver vs1 -path /repository -recurse false

```
The following example recursively creates hashes for all the files and
directories inside the directory "documents":
```

cluster1::> vserver cifs branchcache hash-create -vserver vs1 -path /documents -recurse true

```
[[ID7065c7df3177fe83ba1a94e1a074e9ac]]
= vserver cifs branchcache hash-flush
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Flush all generated BranchCache hashes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs branchcache hash-flush` command deletes all hash data
from the configured hash store.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver whose hash data is to be
deleted.

== Examples

The following example flushes all the hashes for Vserver vs1:
```

cluster1::> vserver cifs branchcache hash-flush -vserver vs1

```
[[IDebc5ac508f408309e1eaae7aec013025]]
= vserver cifs branchcache modify
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the CIFS BranchCache service settings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs branchcache modify` command modifies the configuration
for computing and retrieving BranchCache hash data.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver whose BranchCache
service is to be modified.

``[-versions {v1-enable|v2-enable|enable-all}]`` - Supported BranchCache
Versions::
This optional parameter specifies a list of versions of the BranchCache
protocol that the storage system supports. The default is `enable-all` .
This list can include one or more of the following:
+
  * v1-enable - This option enables BranchCache Version 1.
  * v2-enable - This option enables BranchCache Version 2.
  * enable-all - This option enables all supported versions of
BranchCache.



``[-operating-mode <BranchCache Mode>]`` - CIFS BranchCache Operating
Modes::
This optional parameter specifies the mode in which the BranchCache
service operates. The default is `per-share` . Possible values include:
+
  * disable - This option disables the BranchCache service for the
Vserver.
  * all-shares - This option enables the BranchCache service for all the
shares on this Vserver.
  * per-share - This option enables the BranchCache service on a per-share
basis. You can enable the BranchCache service on an existing share by
adding the ``_branchcache_`` flag in the `-share-properties` parameter of
the xref:{relative_path}vserver-cifs-share-modify.html[vserver cifs share
```

```
modify] command.



``[-hash-store-max-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Size of
the Hash Store::
This optional parameter specifies the maximum size to use for the hash
data. If the size of the hash data exceeds this value, older hashes are
deleted to make room for newer hashes. The default is 1 GB.

``[-flush-hashes {true|false}]`` - Delete Existing Hashes::
This parameter specifies whether to keep or delete all the existing
hashes. This must be set to true when modifying the server key.

``[-hash-store-path <text>]`` - Path to Hash Store::
This parameter specifies an existing directory into which the hash data is
stored. Read-only paths, such as snapshot directories, are not allowed.

``[-server-key <text>]`` - Encryption Key Used to Secure the Hashes::
This optional parameter specifies a server key that the BranchCache
service uses to prevent clients from impersonating the BranchCache server.
If you specify this parameter, all existing hashes for the Vserver are
deleted.

== Examples

The following example modifies the BranchCache service on the Vserver
named vs1. The path to the hash store is /vs1_hash_store_2, the server key
used to secure the hashes is set to "new vserver secret", all existing
hashes are removed, the service supports all BranchCache versions, and is
enabled on a per-share basis.
```

cluster1::> vserver cifs branchcache modify -vserver vs1 -hash-store-path /vs1_hash_store_2 -server-key "new vserver secret" -flush-hashes true -versions enable-all -operating-mode per-share

```
 The following example modifies the BranchCache service on the Vserver
vs1. The service is enabled on all the shares of the Vserver, supports
BranchCache version 1, and supports a maximum of 1 TB of BranchCache
hashes.
```

cluster1::> vserver cifs branchcache modify -vserver vs1 -operating-mode all-shares -versions v1-enable -hash -store-max-size 1TB

```
== Related Links
```

```
* xref:{relative_path}vserver-cifs-share-modify.html[vserver cifs share
modify]




[[IDfc70daa8e075f508509f9a4a2b218cc2]]
= vserver cifs branchcache show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the CIFS BranchCache service status and settings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs branchcache show` command displays information about the
BranchCache configuration for the Vserver. The command output depends on
the parameter or parameters specified with the command. If you do not
specify any parameters, the command displays the following information:

* Operating Mode
* Allowed Versions
* Maximum Size
* Path

You can specify additional parameters to display only information that
matches those parameters.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command
displays only the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information for the
specified Vserver.
```

```
``[-versions {v1-enable|v2-enable|enable-all}]`` - Supported BranchCache
Versions::
If you specify this parameter, the command displays information for the
Vservers that support the specified BranchCache versions.

``[-hash-store-path <text>]`` - Path to Hash Store::
If you specify this parameter, the command displays information for
Vservers that store their hashes at the specified location.

``[-hash-store-max-size {<integer>[KB|MB|GB|TB|PB]}]`` - Maximum Size of
the Hash Store::
If you specify this parameter, the command displays information for
Vservers that have a maximum hash store size that is set to the specified
value.

``[-server-key <text>]`` - Encryption Key Used to Secure the Hashes::
If you specify this parameter, the command displays information for
Vservers that have the specified server key.

``[-operating-mode <BranchCache Mode>]`` - CIFS BranchCache Operating
Modes::
If you specify this parameter, the command displays information for
Vservers whose BranchCache configuration operates in the specified mode.

== Examples

The following example displays a subset of the information about the
BranchCache service in the cluster.
```

cluster1::> vserver cifs branchcache show
Operating Allowed Max
Vserver Mode Versions Size Path
--- --- --- --- ---
vs1 per_share enable_all 1GB /hash_dir/

```
The following example displays all information about all the Vservers with
BranchCache configurations.
```

cluster1::> vserver cifs show -instance
Vserver: vs1
Supported Versions of BranchCache: enable_all
Path to Hash Store: /hash_dir/
Maximum Size of the Hash Store: 1GB
Encryption Key Used to Secure the Hashes: asdad
CIFS BranchCache Operating Modes: per_share

```
 The following example displays information about BranchCache
configurations that store the hash data at the location
/branchcache_hash_store.
```

cluster1::> vserver cifs branchcache show -hash-store-path /branchcache_hash_store
Operating Allowed Max
Vserver Mode Versions Size Path
---------- --------- ---------- ------ ----------
vs1 per_share enable_all 1GB /branchcache_hash_store

```
[[IDadf8684b253330fa58df44691c666b7d]]
= vserver cifs cache name-to-sid delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache name-to-sid delete-all` command removes all of the
Windows user cache entries cached by the Windows name for the specified
Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the name-to-sid cache
entries need to be deleted.

== Examples

The following example shows how to delete all of the cached name-to-sid
entries for Vserver vs0:
```

cluster1::> vserver cifs cache name-to-sid delete-all -vserver vs0

```
[[ID3cbdc00e0f9c83eb2adea6f0c8077770]]
= vserver cifs cache name-to-sid delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache name-to-sid delete` command removes the Windows
user cache entries cached by the Windows name. If cache propagation is
enabled, the corresponding sid-to-name cache entry will also be removed.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the name-to-sid cache
entry needs to be deleted.

``-win-name <text>`` - Windows Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Windows name for which the cached entry
needs to be deleted.

== Examples

The following example shows how to delete the name-to-sid cache entry for
Vserver vs0 with Windows name user1:
```

cluster1::> vserver cifs cache name-to-sid delete -vserver vs0 -win-name user1

```
[[ID8fdf42429704ba5432c7b6518946de59]]
= vserver cifs cache name-to-sid show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display name-to-sid cache entries

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache name-to-sid show` command displays the Windows
user information cached by Windows name.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the Windows user
entries cached by the Windows name.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the Windows user
entries that are cached by the Windows name need to be displayed.

``[-win-name <text>]`` - Windows Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Windows name for which the cached
entries need to be displayed.

``[-sid <text>]`` - SID [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached Windows
user entries that have the specified security identifier (SID).

``[-sid-type <integer>]`` - SID type [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached Windows
user entries that have the specified security identifier (SID) type.

``[-flags <integer>]`` - Flags [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows user
entries cached by the Windows name that have the specified flags.

``[-domain-name <text>]`` - Domain Name [.advaced-priv]#(privilege:
advanced)#::
```

```
Use this parameter to display information only about the Windows user
entries cached by the Windows name that have the specified domain name.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows user
entries that were cached at the specified time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the user entries
cached by the Windows name that have the specified lookup source.

== Examples

The following example shows how to display all of the Windows users which
are cached by the Windows name:
```

cluster1::> vserver cifs cache name-to-sid show

```
The following example shows how to display all of the Windows user entries
cached by the Windows name for Vserver vs0:
```

cluster1::> vserver cifs cache name-to-sid show -vserver vs0

```
[[IDcfd9d9fcda2cbbd8dcd495825790f978]]
= vserver cifs cache settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify CIFS Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache settings modify` command modifies the Windows
users cache configuration of the specified Vserver.

== Parameters
```

```
``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the Windows users
cache settings need to be modified.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
Windows users database. The value ``_true_`` means the cache is enabled
and the value ``_false_`` means the cache is disabled. The default value
for this parameter is ``_false_`` .

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
negative entries. Negative entries means the entries which are not present
in the Windows users database and the look-up fails. The default value for
this parameter is ``_true_`` . Negative cache is disabled by default if
the parameter ``_is-enabled_`` is set to ``_false_`` .

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes, and seconds)
for which the positive entries need to be cached. The positive entries
means the entries which are present in the Windows users database and the
look-up succeeds. The default value is 24 hours.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes, and seconds)
for which the negative entries need to be cached. The default value is 5
minutes.

``[-is-propagation-enabled {true|false}]`` - Is Propagation Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify whether the cached user entries need to be
propagated to the sid-to-name cache. The default value is ``_true_`` .
Specify ``_false_`` to disable propagation.

== Examples

The following example shows how to modify the Windows users cache
configuration settings for Vserver vs0:
```

cluster1::> vserver cifs cache settings modify -vserver vs0 -ttl 600 -negative-ttl 300

```
The following example shows how to disable the Windows users cache for
Vserver vs0:
```

cluster1::> vserver cifs cache settings modify -vserver vs0 -is-enabled false

```
[[ID179e64bfa2e50854e7d366accdbf8e84]]
= vserver cifs cache settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display CIFS Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache settings show` command displays information about
the Windows users cache configuration of the specified Vserver.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the Windows users
cache configuration settings.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information about the Windows users cache
configuration settings for the Vserver you specify.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows users
cache configuration settings that have the specified cache enabled
setting. Value ``_true_`` displays only the entries that have cache
enabled and value ``_false_`` displays only the entries that have cache
```

```
disabled.

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows users
cache configuration settings that have the specified negative cache
enabled setting. Value ``_true_`` displays only the entries that have
negative cache enabled and value ``_false_`` displays only the entries
that have negative cache disabled.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows users
cache configuration settings that have the specified Time to Live.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows users
cache configuration settings that have the specified negative Time to
Live.

``[-is-propagation-enabled {true|false}]`` - Is Propagation Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows users
cache configuration settings that have the specified propagation enabled
setting. Value ``_true_`` displays only the entries that have the
propagation of cached entries to sid-to-name cache enabled and value
``_false_`` displays only the entries that have the propagation of cached
entries to sid-to-name cache disabled.

== Examples

The following example shows how to display the Windows users cache
configuration settings for all the Vservers:
```

cluster1::> vserver cifs cache settings show

```
The following example shows how to display the Windows users cache
configuration settings for Vserver vs0:
```

cluster1::> vserver cifs cache settings show -vserver vs0

```
The following example shows how to display the Windows users cache
configuration settings that have cache disabled:
```

cluster1::> vserver cifs cache settings show -is-enabled false

```
[[IDd21c0ceb3ab002a08a04553eaa7fe7cd]]
= vserver cifs cache sid-to-name delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache sid-to-name delete-all` command removes all of the
Windows user cache entries cached by the security identifier (SID) for the
specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the sid-to-name cache
entries need to be deleted.

== Examples

The following example shows how to delete all the cached sid-to-name
entries for Vserver vs0:
```

cluster1::> vserver cifs cache sid-to-name delete-all -vserver vs0

```
[[IDdb772a7b29d67c51c1850634b7f08fbd]]
= vserver cifs cache sid-to-name delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache sid-to-name delete` command removes the Windows
user cache entries cached by security identifier (SID). If cache
propagation is enabled, the corresponding name-to-sid cache entry will
also be removed.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the sid-to-name cache
entry needs to be deleted.

``-sid <text>`` - SID [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the security identifier (SID) for which the
cached entry needs to be deleted.

== Examples

The following example shows how to delete the sid-to-name cache entry for
Vserver vs0 with SID S-1-5-21-1380078113-1824080971-954447143-1152:
```

cluster1::> vserver cifs cache sid-to-name delete -vserver vs0 -sid S-1-5-21-1380078113-1824080971-954447143-1152

```
[[ID3eaded5b4911aa807a986b973b942cec]]
= vserver cifs cache sid-to-name show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display sid-to-name cache entries

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs cache sid-to-name show` command displays the Windows
user information cached by security identifier (SID).

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the Windows user
entries cached by the security identifier (SID).

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the Windows user
entries that are cached by the security identifier (SID) need to be
displayed.

``[-sid <text>]`` - SID [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached Windows
user entries that have the specified security identifier (SID).

``[-win-name <text>]`` - Windows Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Windows name for which the cached
entries need to be displayed.

``[-sid-type <integer>]`` - SID type [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached Windows
user entries that have the specified security identifier (SID) type.

``[-sid-mode <integer>]`` - SID mode [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached Windows
user entries that have the specified security identifier (SID) mode.
```

```
``[-flags <integer>]`` - Flags [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows user
entries cached by the security identifier (SID) that have the specified
flags.

``[-domain-name <text>]`` - Domain Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the Windows user
entries cached by the security identifier (SID) that have the specified
domain name.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows user
entries that were cached at the specified time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the Windows user
entries cached by the security identifier (SID) that have the specified
lookup source.

== Examples

The following example shows how to display all of the Windows users which
are cached by the security identifier (SID):
```

cluster1::> vserver cifs cache sid-to-name show

```
The following example shows how to display all of the Windows user entries
cached by the security identifier (SID) for Vserver vs0:
```

cluster1::> vserver cifs cache sid-to-name show -vserver vs0

```
[[ID34eccfe56eb8357323c215fb00cf62a0]]
= vserver cifs character-mapping create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create character mapping on a volume
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs character-mapping create` command creates the CIFS
character mapping for the specified volume on a particular Vserver.

NOTE: Choose target characters in the "Private Use Area" of Unicode in the
following range: U+E000...U+F8FF.



NOTE: The target Unicode characters must not appear in existing file
names; otherwise, unwanted character mappings would occur, resulting in
clients being unable to access mapped files. For example, if ":" is mapped
to "-" but "-" appears in files normally, a Windows client using the
mapped share to access a file named "a-b" would have its request mapped to
the NFS name "a:b", which is not the desired file.



The `vserver cifs character-mapping create` command is not supported for
FlexGroup volumes.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which a volume is located for
which you are creating the character mapping. If only one data Vserver
exists, you do not need to specify this parameter.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume for which you are creating
the character mapping.

``-mapping <text>,...`` - Character Mapping::
This parameter specifies the mapping of the invalid CIFS filename
characters to valid CIFS filename characters. The mapping consists of a
list of source-target character pairs separated by ":". The characters are
Unicode characters entered using hexadecimal digits. For example: 3C:E03C.
+
NOTE: The permissible Unicode character set for source mapping is: 0x01-
0x19, 0x5C, 0x3A, 0x2A, 0x3F, 0x22, 0x3C, 0x3E, 0x7C, 0xB1.
+
```

```
== Examples

The following example creates a character mapping for a volume vol1 on
Vserver vs1.
```

cluster1::> vserver cifs character-mapping create -volume vol1 -mapping 3c:e17c, 3e:f17d, 2a:f745
cluster1::> vserver cifs character-mapping show

Vserver Volume Name Character Mapping
--- --- ---
vs1 vol1 3c:e17c, 3e:f17d, 2a:f745

```
[[ID182c54efb22d9e6ebc7e63ce0f34b7fd]]
= vserver cifs character-mapping delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete character mapping on a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs character-mapping delete` command deletes the CIFS
character mapping for the specified volume on a particular Vserver.

The `vserver cifs character-mapping delete` command is not supported for
FlexGroup volumes.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which a Volume is located for
which you are deleting the character mapping. If only one data Vserver
exists, you do not need to specify this parameter.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume for which you are deleting
the character mapping.

== Examples

The following example deletes all character mappings for a volume vol1 on
Vserver vs1.
```

cluster1::> vserver cifs character-mapping delete -volume vol1

```
[[ID036242f627c75c968e0f119514ae75b1]]
= vserver cifs character-mapping modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Modify character mapping on a volume

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs character-mapping modify` command modifies the CIFS
character mapping for the specified volume on a particular Vserver.

You can modify a particular volume's character mapping by specifying the
following two parameters in the modify command:

* Vserver associated with the volume
* Name of the Volume



NOTE: Choose target characters in the "Private Use Area" of Unicode in the
following range: U+E000...U+F8FF.



NOTE: The target Unicode characters must not appear in existing file
names; otherwise, unwanted character mappings would occur, resulting in
clients being unable to access mapped files. For example, if ":" is mapped
to "-" but "-" appears in files normally, a Windows client using the
mapped share to access a file named "a-b" would have its request mapped to
the NFS name "a:b", which is not the desired file.



The `vserver cifs character-mapping modify` command is not supported for
FlexGroup volumes.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which a Volume is located for
which you are modifying the character mapping. If only one data Vserver
exists, you do not need to specify this parameter.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume for which you are
modifying the character mapping.
```

```
``[-mapping <text>,...]`` - Character Mapping::
This parameter specifies the mapping of the invalid CIFS filename
characters to valid CIFS filename characters. The mapping consists of a
list of source-target character pairs separated by ":". The characters are
Unicode characters entered using hexadecimal digits. For example: 3C:E03C.
+
NOTE: The permissible Unicode character set for source mapping is: 0x01-
0x19, 0x5C, 0x3A, 0x2A, 0x3F, 0x22, 0x3C, 0x3E, 0x7C, 0xB1.
+


== Examples

The following example modifies a character mapping for a volume vol1 on
Vserver vs1.
```

cluster1::> vserver cifs character-mapping modify -volume vol1 -mapping 3c:e17d, 3e:f17e, 2a:f746
cluster1::> vserver cifs character-mapping show

Vserver Volume Name Character Mapping
--- --- ---
vs1 vol1 3c:e17d, 3e:f17e, 2a:f746

```
[[IDae42a39ee2ce8ac1779ac20e5efece16]]
= vserver cifs character-mapping show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display character mapping on volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs character-mapping show` command displays information
about character mapping configured for volumes. The command output depends
on the parameter or parameters specified with the command. If you do not
specify any parameters, the command displays the following information
about character mapping configured for volumes:

* Vserver name
```

```
* Volume name
* Character mapping



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify this parameter, the command displays only the fields that
you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information about
character mapping configured for all the volumes that belong to the
specified Vserver.

``[-volume <volume name>]`` - Volume Name::
If you specify this parameter, the command displays information about the
character mapping configured for all the volumes that match the specified
volume name.

``[-mapping <text>,...]`` - Character Mapping::
If you specify this parameter, the command displays information about the
character mapping configured for all volumes that match the specified
mapping.

== Examples

The following example displays information about all character mappings
configured for volumes
```

cluster1::> vserver cifs character-mapping show

Vserver Volume Name Character Mapping
--- --- ---
vs1 vol1 3c:e17d, 3e:f17e

```
[[ID75bcfcbc16d577735a5f65915e4869d8]]
= vserver cifs connection show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays established CIFS connections

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs connection show` command displays information about
established CIFS connections.

== Parameters

{ ``[-fields <fieldname>,...]``::
Use this parameter to display only the specified fields

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
Use this parameter to display information about CIFS connections on the
specified node.

``[-vserver <vserver name>]`` - Vserver::
Use this parameter to display information about CIFS connections on the
specified CIFS-enabled SVM.

``[-connection-id <integer>]`` - Connection ID::
Use this parameter to display information about CIFS connections that
match the specified connection ID.

``[-session-id <integer>,...]`` - Session ID::
Use this parameter to display information about CIFS connections that
match the specified session ID.

``[-workstation-ip <IP Address>]`` - Workstation IP Address::
Use this parameter to display information about CIFS connections that are
established through the specified data LIF IP address.

``[-workstation-port <integer>]`` - Workstation Port Number::
Use this parameter to display information about CIFS connections that are
opened from the specified Port number.

``[-lif-ip <IP Address>]`` - Incoming Data LIF IP Address::
```

```
Use this parameter to display information about CIFS connections that are
opened from the specified IP address.

``[-network-context-id <integer>]`` - Network Context ID [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information about CIFS connections that
match the specified network context ID.

== Examples

The following example displays information about all CIFS connections:
```

cluster1::> vserver cifs connection show
Node: node1
Vserver: vs1
Connection Session Workstation
ID IDs Workstation IP Port LIF IP
---------- ---------- ---------- ---------- ----------
127834 1,2 172.17.193.172 15536 10.53.50.42

```
The following example displays information about a CIFS connection at
advanced privilege level:
```

cluster1::*> vserver cifs connection show
Node: node1
Vserver: vs1
Connection Session Workstation Network
ID IDs Workstation IP Port LIF IP Context ID
---------- ---------- ---------- ---------- ---------- ----------
127834 1,2 172.17.193.172 15536 10.53.50.42 2

```
The following example displays information about a CIFS connection with
session-id 1:
```

cluster1::*> vserver cifs connection show -session-id 1 -instance

Vserver: vs1
Node: node1
Connection ID: 127834
Session ID: 1
Workstation IP Address: 172.17.193.172
Workstation Port Number: 15536
Incoming Data LIF IP Address: 10.53.50.42
Network Context ID: 2

```
[[ID5b8306a17766e5c3fc01c92ab90c464c]]
= vserver cifs domain discovered-servers reset-servers
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset and rediscover servers for a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain discovered-servers reset-servers` command
discards information the storage system has stored about domain
controllers, LDAP, and NIS servers. After that, it begins the discovery
process to reacquire current information about external servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver.

== Examples

The following is an example use of this command. It produces no output.
```

cluster1::> vserver cifs domain discovered-servers reset-servers

cluster1::>

```
[[IDb3e5f4d7ba8d71c68698eff584c3014c]]
= vserver cifs domain discovered-servers show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display discovered server information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver cifs domain discovered-servers show` command displays
information about the discovered servers for the CIFS domains of one or
more Vservers. Server displays are grouped by node and Vserver, and each
group is preceded by the node and Vserver identification. Within each
grouping, the server display is limited to those associated with the
domain specified by the domain parameter, if it is present.

== 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>|local}]`` - Node::
If you use this parameter, the command only displays servers for the
specified node.

``[-vserver <vserver name>]`` - Vserver::
If you use this parameter, the command only displays servers for the
specified Vserver.

``[-domain <TextNoCase>]`` - Fully Qualified Domain Name::
If you use this parameter, the command only displays servers in the
specified domain.

``[-type {Unknown|KERBEROS|MS-LDAP|MS-DC|LDAP}]`` - Server Type::
If you use this parameter, the command only displays servers of the
specified type.

``[-name <text>]`` - Server Name::
If you use this parameter, the command only displays servers the with the
specified name. This can result in multiple lines because the same server
may provide multiple services.

``[-address <InetAddress>]`` - Server Address::
If you use this parameter, the command only displays servers with the
specified IP address. This can result in multiple lines because the same
server may provide multiple services.
```

```
``[-preference {unknown|preferred|favored|adequate}]`` - Preference::
If you use this parameter, the command only displays servers of the
specified preference level.

``[-status {OK|unavailable|slow|expired|undetermined|unreachable}]`` -
Status::
If you use this parameter, the command only displays servers of the
specified status.

``[-dc-functional-level
{win2000|unknown|win2003|win2008|win2008r2|win2012|win2012r2|win2016|winth
reshold}]`` - DC Functional Level::
If you use this parameter, the command only displays servers with the
specified functional level.

``[-is-dc-read-only {true|false}]`` - Is DC Read Only::
If this parameter is set to true, the command only displays servers with
read only domain controller. If set to false, the command only displays
servers with writable domain controller.

== Examples

The following example display shows the information provided by this
command.
```

cluster1::> vserver cifs domain discovered-servers show

Node: node1
Vserver: vs1

Domain Name Type Preference DC-Name DC-Address Status
--- --- --- --- --- ---
"" NIS preferred 192.168.10.222 192.168.10.222 OK
example.com MS-LDAP adequate DC-1 192.168.192.24 OK
example.com MS-LDAP adequate DC-2 192.168.192.25 OK
example.com MS-DC adequate DC-1 192.168.192.24 OK
example.com MS-DC adequate DC-2 192.168.192.25 OK
5 entries were displayed.

```
[[ID653b0e0fc4bcf77d506073debe2a5688]]
= vserver cifs domain discovered-servers discovery-mode modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify Server Discovery Mode

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs domain discovered-servers discovery-mode modify` command
modifies the configuration for the server discovery mode of one or more
Data Vservers. This option controls the way domain controllers(DCs) are
discovered.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which you want to modify the
server discovery mode.

``[-mode {all|site|none}]`` - Server Discovery Mode [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify the server discovery mode for the Vserver.
Following are the possible values for this parameter:
+
  *  all - Discover all the DCs in the domain, including the DCs local to
the site. If trusted domains are present, then discover all the KDCs in
the trusted domain as well, to be used for kerberos communication
(default).
  *  site - Discover the DCs local to the site. If trusted domains are
present, then discover KDCs local to the trusted domain site as well, to
be used for kerberos communication.
  *  none - Discover nothing. Depend only on preferred-dc configured.



== Examples

The following example disables server discovery for a Vserver.
```

cluster1::*> vserver cifs domain discovered-servers discovery-mode modify -vserver vs1 -mode none

```
[[ID6b3592f7636c114a22209b7e7962e9d7]]
= vserver cifs domain discovered-servers discovery-mode show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Server Discovery Mode

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs domain discovered-servers discovery-mode show` command
displays information about the discovery mode for domain controllers(DCs)
of one or more Vservers.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you use this parameter, the command only displays discovery mode for
the specified Vserver.

``[-mode {all|site|none}]`` - Server Discovery Mode [.advaced-
priv]#(privilege: advanced)#::
If you use this parameter, the command only displays Vservers with the
specified mode.

== Examples

The following example shows the server discovery mode for all Vservers.
```

cluster1::*> vserver cifs domain discovered-servers discovery-mode show
Vserver Mode
--- ---
vs1 all
vs2 site
vs3 none
3 entries were displayed.

```
[[IDc7c0e38b094432cb9b29ea151f635a84]]
= vserver cifs domain name-mapping-search add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to the list of trusted domains for name-mapping

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain name-mapping-search add` command adds one or more
trusted domains to the list of trusted domains to be used in preference to
all others by the specified Vserver for looking up Windows user names when
performing UNIX user to Windows user name-mapping. If a list already
exists for the specified vserver, the new list is merged with the existing
list. This command is not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which you want to add
trusted domains.

``-trusted-domains <domain name>,...`` - Trusted Domains::
This parameter specifies a comma-delimited list of fully-qualified domain
names of the trusted domains for the home domain.

== Examples

The following example adds two trusted domains (cifs1.example.com and
cifs2.example.com) to the preferred list used by Vserver vs1:
```

cluster1::> vserver cifs domain name-mapping-search add -vserver vs1 -trusted-domains cifs1.example.com, cifs2.example.com

```
[[IDe1c1e28540de76d4b59dee7a896fe81e]]
= vserver cifs domain name-mapping-search modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the list of trusted domains for name-mapping search

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain name-mapping-search modify` command modifies the
current list of trusted domains to be used in preference to all others by
the specified Vserver to lookup Windows user names when performing UNIX
user to Windows user name-mapping. The new list overwrites the existing
list. This command is not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which you want to
modify the trusted domain list.

``-trusted-domains <domain name>,...`` - Trusted Domains::
This parameter specifies a comma-delimited list of fully qualified domain
names of the trusted domains of the home domain.

== Examples

The following example modifies the trusted domain list used by Vserver
vs1:
```

cluster1::> vserver cifs domain name-mapping-search modify -vserver vs1 -trusted-domains cifs3.example.com

```
[[ID0b9b66e867111191cb7ceb660dfe7e2e]]
= vserver cifs domain name-mapping-search remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from the list of trusted domains for name-mapping search

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain name-mapping-search remove` command removes one
or more trusted domains from the list used by the specified Vserver to
lookup Windows user names when performing UNIX user to Windows user name-
mapping. If a list of trusted domains is not provided, the entire trusted
domain list for the specified Vserver is removed. This command is not
supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver from which you want to
remove trusted domains.

``[-trusted-domains <domain name>,...]`` - Trusted Domains::
This parameter specifies a comma-delimited list of trusted domains of the
home domain.

== Examples

The following example removes two trusted domains from the list used by
Vserver vs1:
```

cluster1::> vserver cifs domain name-mapping-search remove -trusted-domains cifs1.example.com, cifs2.example.com

```
[[IDc6b50971d46c2a76ff03f0d5a6192cdf]]
= vserver cifs domain name-mapping-search show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of trusted domains for name-mapping searches

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain name-mapping-search show` command displays
information about trusted domains of the home domain by Vserver.

== 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.

``[-vserver <vserver name>]`` - Vserver::
This parameter specifies the name of the Vserver for which you want to
display information about the trusted domains.

``[-trusted-domains <domain name>,...]`` - Trusted domains::
This parameter specifies a comma-delimited list of fully qualified domain
names of trusted domains for which you want to display information.

== Examples

The following example displays information about all preferred trusted
domains:
```

cluster1::> vserver cifs domain name-mapping-search show
Vserver Trusted Domains
---------- ----------
vserver_1 CIFS1.EXAMPLE.COM

```
[[ID975c2606b1e36d6a1cb403b579b59423]]
= vserver cifs domain password change
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Generate a new password for the CIFS server's machine account and change
it in the Windows Active Directory domain.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain password change` changes the domain account
password for a CIFS server. This command is not supported for workgroup
CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for whose CIFS server you
want to change the domain account password.

== Examples

The following example changes the password for the CIFS server on a
Vserver named vs1.
```

cluster1::> vserver cifs domain password change -vserver vs1

cluster1::>

```
[[IDc926b49c9def120dc6c2c1dff8757127]]
= vserver cifs domain password reset
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset the CIFS server's machine account password in the Windows Active
Directory domain.
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain password reset` command resets the domain account
password for a CIFS server. This may be required if the password stored
along with the machine account in the Windows Active Directory domain is
changed or reset without the Vserver's knowledge. The operation requires
the credentials for a user with permission to reset the password in the
organizational unit (OU) that the machine account is a member of. This
command is not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for whose CIFS server you
want to reset the domain account password.

``[-client-id <text>]`` - Application (Client) ID of Deployed Azure
Application::
This parameter specifies the Application (Client) ID of Deployed "Entra
ID" Application.

``[-tenant-id <text>]`` - Directory (Tenant) ID of Deployed Azure
Application::
This parameter specifies the Directory (Tenant) ID of Deployed "Entra ID"
Application.

``[-client-certificate <text>]`` - Certificate for Azure Application::
This parameter specifies the Certificate to authenticate with Deployed
"Entra ID" Application.

== Examples

The following example resets the password for the CIFS server on a Vserver
named vs1.
```

cluster1::> vserver cifs domain password reset -vserver vs1

Enter your user ID: Administrator
Enter your password:

cluster1::>

```
[[IDc079c380d193841373d3cc89e35ea71a]]
= vserver cifs domain password schedule modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the domain account password change schedule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain password schedule modify` command enables you to
modify a domain account password change schedule for a CIFS server. This
command is not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver containing the CIFS server for
which you want to change the domain account password.

``[-is-schedule-enabled {true|false}]`` - Is Password Change Schedule
Enabled::
This specifies whether the domain account password change schedule is
enabled.

``[-schedule-weekly-interval <integer>]`` - Interval in Weeks for Password
Change Schedule::
This specifies the number of weeks after which the scheduled domain
account password change must occur.

``[-schedule-randomized-minute <integer>]`` - Minutes Within Which
Schedule Start Can be Randomized::
This specifies the minutes within which the scheduled domain account
password change must begin.

``[-schedule-day-of-week <cron_dayofweek>]`` - Day of Week for Password
Change Schedule::
This sets the day of week when the scheduled domain account password
change occurs.

``[-schedule-time-of-day <HH:MM:SS>]`` - Start Time for Password Change
Schedule::
```

```
This sets the time in HH:MM:SS at which the scheduled domain account
password change starts.

== Examples

The following example enables the domain account password change schedule
and configures it to run at any time between 23:00:00 to 00:59:00 (one
hour before midnight to one hour after midnight) on every 4th Sunday.
```

cluster1::> vserver cifs domain password schedule modify -is-schedule-enabled true -schedule-randomized -minute 120 -schedule-weekly-interval 4 -schedule-time-of-day 23:00:00 -schedule-day-of-week sunday

```
[[ID509a8f735994cb55eaebed48821c9a78]]
= vserver cifs domain password schedule show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the domain account password change schedule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain password schedule show` command displays the
domain account password change schedule configuration. It displays the
following fields:

* Vserver: Vserver for which the schedule is configured
* Schedule Enabled: Whether the schedule is enabled or disabled for this
Vserver
* Schedule Interval: Weeks after which the password change schedule occurs
again for this Vserver
* Schedule Randomized Within: Minutes within which the schedule must begin
for this Vserver
* Schedule: Password change schedule currently set on this Vserver
* Last Successful Password Change/Reset Time: Time at which the last
password change or reset happened successfully on this Vserver
* Warning: Warning message, applicable only when the change password job
is deleted with the feature still enabled on this Vserver
```

```
== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information for the
specified Vserver.

``[-is-schedule-enabled {true|false}]`` - Is Password Change Schedule
Enabled::
If you specify this parameter, the command displays information for all
the Vservers on which the is-schedule-enabled value applies.

``[-schedule-weekly-interval <integer>]`` - Interval in Weeks for Password
Change Schedule::
If you specify this parameter, the command displays information for all
the Vservers on which the schedule-weekly-interval value applies.

``[-schedule-randomized-minute <integer>]`` - Minutes Within Which
Schedule Start Can be Randomized::
If you specify this parameter, the command displays information for all
the Vservers on which the schedule-randomized-minute value applies.

``[-schedule-last-changed <text>]`` - Last Successful Password
Change/Reset Time::
If you specify this parameter, the command displays information for all
the Vservers on which the schedule-last-changed value applies.

``[-schedule-description <text>]`` - Schedule Description::
If you specify this parameter, the command displays information for all
the Vservers on which the schedule-description value applies.

``[-schedule-warn-msg <text>]`` - Warning Message in Case Job Is Deleted::
If you specify this parameter, the command displays information for all
the Vservers on which the schedule-warn-msg value applies.

== Examples

The following example shows the domain account password change schedule
configuration when the password change feature is enabled for Vserver vs1.
```

```
cluster1::> vserver cifs domain password schedule show
Vserver: vs1
Schedule Enabled: true
Schedule Interval: 4 week
Schedule Randomized Within: 120 min
Schedule: Fri@23:00
Last Changed At: Thu Apr 4 02:35:23 2013
```

```
The following example shows the domain account password change schedule
configuration when the password change job has been accidentally deleted.
```

```
cluster1::> vserver cifs domain password schedule show
Vserver: vs1
Schedule Enabled: true
Schedule Interval: 4 week
Schedule Randomized Within: 120 min
Schedule: Fri@23:00
Last Changed At: Thu Apr 4 02:35:23 2013
Warning: Password change job was deleted. Re-enable the password change schedule.
```

```
[[ID4f2b5ec786abbf1ae2d20a5fb4fb87e1]]
= vserver cifs domain preferred-dc add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to a list of preferred domain controllers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain preferred-dc add` command adds one or more domain
controllers to be used in preference to all others by the specified
Vserver for interactions with the specified domain. If a list already
exists for the specified domain, the new list is merged with the existing
list. This command is not supported for workgroup CIFS servers.

NOTE: Each Vserver discovers domain controllers and attempts to sort them
internally based on real-world performance. Therefore it should not be
necessary to create a preferred list of domain controllers under most
circumstances.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which you want to add
preferred domain controllers.

``-domain <TextNoCase>`` - Fully Qualified Domain Name::
This parameter specifies the fully-qualified name of the domain that the
domain controllers belong to.

``-preferred-dc <InetAddress>,...`` - Preferred Domain Controllers::
This parameter specifies a comma-delimited list of IP addresses for domain
controllers that belong to the domain specified in the `-domain`
parameter.

``[-skip-config-validation {true|false}]`` - Skip Configuration
Validation::
Use this parameter to skip the Preferred-DC configuration validation.
+
The hosts specified with the `-DC-servers` parameter are validated to
verify that each of the DC servers are reachable, and is providing
NETLOGON services.
+
The validation fails if there is no valid Preferred-DC server.

== Examples

The following example adds two domain controllers (192.168.0.100 and
192.168.0.101) to the preferred list used by Vserver vs1 when connecting
to the example.com domain:
```

cluster1::> vserver cifs domain preferred-dc add -vserver vs1 -domain example.com -preferred-dc 192.168.0.100,192.168.0.101

```
[[ID74e4631ce78d50efddb4f033c3fe5ea1]]
= vserver cifs domain preferred-dc check
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display validation status of the Preferred-DC configuration

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

Use the `vserver cifs domain preferred-dc check` command to check the
status of configured preferred DC on a particular vserver.

== 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.

``-vserver <vserver>`` - Vserver::
Use this parameter to specify the Vserver whose preferred DC needs to be
validated.

``[-server-ip <text>,...]`` - DC Address::
Use this parameter to display IP-address of the configured CIFS Preferred-
DC servers.

``[-domain <TextNoCase>,...]`` - Domain Name::
Use this parameter to display Domain name of the configured CIFS
Preferred-DCs.

``[-status {down|up}]`` - DC Status::
Use this parameter to display information only about CIFS Preferred-DC
servers with a status that matches the value you specify.

``[-status-details <text>,...]`` - Status Details::
Use this parameter to display information only about CIFS Preferred-DC
servers with status details that match the value you specify.

== Examples

The following example checks the connectivity of preferred DC on vserver
vs0.
```

cluster1::> vserver cifs domain preferred-dc check -vserver vs0

Vserver : vs0

```
Domain Name DC Address Status Status Details
----------- ----------- ------ --------------
example.com 1.1.1.1 up Response time (msec): 426
example.com 1.1.1.2 up Response time (msec): 425
example1.com 2.2.2.2 up Response time (msec): 423
example2.com 3.3.3.3 up Response time (msec): 422
```

```
[[ID8b95e5dfd076630d29f5051d00ae5775]]
= vserver cifs domain preferred-dc remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from a list of preferred domain controllers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain preferred-dc remove` command removes one or more
domain controllers from the list used by the specified Vserver for
interactions with the specified domain. If a list of preferred domain
controllers is not provided, the entire list for the specified domain is
removed. This command is not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver from which you want to
remove preferred domain controllers.

``-domain <TextNoCase>`` - Fully Qualified Domain Name::
This parameter specifies the fully-qualified name of the domain that the
domain controllers belong to.

``[-preferred-dc <InetAddress>,...]`` - Preferred Domain Controllers::
This parameter specifies a comma-delimited list of IP addresses for domain
controllers that belong to the domain specified in the `-domain`
parameter.

== Examples

The following example removes one domain controller (192.168.0.101) from
the preferred list used by Vserver vs1 when connecting to the example.com
domain:
```

cluster1::> vserver cifs domain preferred-dc remove -vserver vs1 -domain example.com -preferred-dc 192.168.0.101

```
[[ID80449790e7b2c370adecc9d950aa70fc]]
= vserver cifs domain preferred-dc show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of preferred domain controllers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain preferred-dc show` command displays lists of
preferred domain controllers by Vserver and domain.

== 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.

``[-vserver <vserver name>]`` - Vserver::
This parameter specifies the name of the Vserver for which you want to
display preferred domain controllers.

``[-domain <TextNoCase>]`` - Fully Qualified Domain Name::
This parameter specifies the fully-qualified name of the domain of the
domain controllers to display.

``[-preferred-dc <InetAddress>,...]`` - Preferred Domain Controllers::
This parameter specifies a comma-delimited list of IP addresses for domain
controllers to display.

== Examples

The following example displays all preferred domain controllers for all
Vservers:
```

cluster1::> vserver cifs domain preferred-dc show
Vserver Domain Name Preferred Domain Controllers
--- --- ---
vs1 example.com 192.168.0.100, 192.168.0.101

```
[[ID941b5eda7adefd83ee5a2b4281dbaed9]]
= vserver cifs domain trusts rediscover
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset and rediscover trusted domains for a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain trusts rediscover` command discards information
the Vserver has stored about trusted domains. After that, it begins the
discovery process to reacquire current information about trusted domains.
This command is not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver.

== Examples

The following example rediscovers trusted domains. It produces no output.
```

cluster1::> vserver cifs domain trusts rediscover

```
[[ID81fc52cbbb52823ee110e1e7466394dc]]
= vserver cifs domain trusts show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display discovered trusted domain information
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs domain trusts show` command displays information about
the trusted domains for the CIFS home domain of one or more Vservers. The
displayed trusted domain information is grouped by node and Vserver, and
each group is preceded by the node and Vserver identification. This
command is not supported for workgroup CIFS servers.

== 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>|local}]`` - Node::
If you use this parameter, the command displays information only about
trusted domains of the home domains for the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you use this parameter, the command displays information only about
trusted domains of the home domain for the specified Vserver.

``[-home-domain <domain name>]`` - Home Domain Name::
If you use this parameter, the command displays information only about
trusted domains of the home domain with the specified name.

``[-trusted-domain <domain name>,...]`` - Trusted Domain Name::
If you use this parameter, the command displays information only about
trusted domains with the specified name.

== Examples

The following example displays information about all the bidirectional
trusted domains for node-01 and vserver_1.
```

cluster1::> vserver cifs domain trusts show -node node-01 -vserver vserver_1
Node: node-01
Vserver: vserver_1

Home Domain Trusted Domain
--- ---
EXAMPLE.COM CIFS1.EXAMPLE.COM,
CIFS2.EXAMPLE.COM

```
[[IDcf1d2ea91f794a2cc1c5be594c46accb]]
= vserver cifs group-policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change group policy configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy modify` command modifies the group policy
configuration of a CIFS server. This command is not supported for
workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver whose group policy configuration you
want to modify.

``[-status {enabled|disabled}]`` - Group Policy Status::
This parameter specifies whether the CIFS-enabled Vserver's group policy
is enabled or disabled.

== Examples

The following example enables the group policy for CIFS-enabled Vserver
vs1.
```

cluster1::> vserver cifs group-policy modify -vserver vs1 -status enabled

```
[[IDa935f5875cac2db5cbc4e3498d010507]]
= vserver cifs group-policy show-applied
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show currently applied group policy setting

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays only group policy
information that has been applied to the Vserver you specify.

``[-gpo-index <integer>]`` - GPO Index::
If you specify this parameter, the command displays only group policy
information at gpo-index.

== Examples

The following example displays all group policy information about all
group policies that have been applied to a Vserver:
```

cluster1::> vserver cifs group-policy show-applied

Vserver: vs1

GPO Name: Default Domain Policy
Level: Domain
Status: enabled
Advanced Audit Settings:
Object Access:
Central Access Policy Staging: failure
Registry Settings:

Refresh Time Interval: 22
Refresh Random Offset: 8
Hash Publication Mode for BranchCache: per-share
Hash Version Support for BranchCache: all-versions
Security Settings:
Event Audit and Event Log:
Audit Logon Events: none
Audit Object Access: success
Log Retention Method: overwrite-as-needed
Max Log Size: 16384
File Security:
/vol1/home
/vol1/dir1
Kerberos:
Max Clock Skew: 5
Max Ticket Age: 10
Max Renew Age: 7
Privilege Rights:
Take Ownership: usr1, usr2
Security Privilege: usr1, usr2
Change Notify: usr1, usr2
Registry Values:
Signing Required: false
Restrict Anonymous:
No enumeration of SAM accounts: true
No enumeration of SAM accounts and shares: false
Restrict anonymous access to shares and named pipes: true
Combined restriction for anonymous user: no-access
Restricted Groups:
gpr1
gpr2
Central Access Policy Settings:
Policies: cap1
cap2
GPO Name: Resultant Set of Policy
Level: RSOP
Advanced Audit Settings:
Object Access:
Central Access Policy Staging: failure
Registry Settings:
Refresh Time Interval: 22
Refresh Random Offset: 8
Hash Publication Mode for BranchCache: per-share
Hash Version Support for BranchCache: all-versions
Security Settings:
Event Audit and Event Log:
Audit Logon Events: none
Audit Object Access: success
Log Retention Method: overwrite-as-needed
Max Log Size: 16384
File Security:
/vol1/home
/vol1/dir1
Kerberos:

Max Clock Skew: 5
Max Ticket Age: 10
Max Renew Age: 7
Privilege Rights:
Take Ownership: usr1, usr2
Security Privilege: usr1, usr2
Change Notify: usr1, usr2
Registry Values:
Signing Required: false
Restrict Anonymous:
No enumeration of SAM accounts: true
No enumeration of SAM accounts and shares: false
Restrict anonymous access to shares and named pipes: true
Combined restriction for anonymous user: no-access
Restricted Groups:
gpr1
gpr2
Central Access Policy Settings:
Policies: cap1
cap2
2 entries were displayed.

```
[[ID1c6b556963a21b6e29cb51ca86664e50]]
= vserver cifs group-policy show-defined
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show applicable group policy settings defined in Active Directory

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays only group policy
information that has been defined in Active Directory for the Vserver that
you specify.

``[-gpo-index <integer>]`` - GPO Index::
If you specify this parameter, the command displays only group policy
information at gpo-index.

== Examples

The following example displays all group policy information for all group
policies that have been defined in Active Directory:
```

cluster1::> vserver cifs group-policy show-defined

Vserver: vs1

GPO Name: Default Domain Policy
Level: Domain
Status: enabled
Advanced Audit Settings:
Object Access:
Central Access Policy Staging: failure

Registry Settings:
Refresh Time Interval: 22
Refresh Random Offset: 8
Hash Publication Mode for BranchCache: per-share
Hash Version Support for BranchCache : version1
Security Settings:
Event Audit and Event Log:
Audit Logon Events: none
Audit Object Access: success
Log Retention Method: overwrite-as-needed
Max Log Size: 16384
File Security:
/vol1/home
/vol1/dir1
Kerberos:
Max Clock Skew: 5
Max Ticket Age: 10
Max Renew Age: 7
Privilege Rights:
Take Ownership: usr1, usr2
Security Privilege: usr1, usr2
Change Notify: usr1, usr2
Registry Values:
Signing Required: false
Restrict Anonymous:
No enumeration of SAM accounts: true
No enumeration of SAM accounts and shares: false
Restrict anonymous access to shares and named pipes: true
Combined restriction for anonymous user: no-access
Restricted Groups:
gpr1
gpr2
Central Access Policy Settings:
Policies: cap1
cap2
GPO Name: Resultant Set of Policy
Status: enabled
Advanced Audit Settings:
Object Access:
Central Access Policy Staging: failure
Registry Settings:
Refresh Time Interval: 22
Refresh Random Offset: 8
Hash Publication for Mode BranchCache: per-share
Hash Version Support for BranchCache: version1
Security Settings:
Event Audit and Event Log:
Audit Logon Events: none
Audit Object Access: success
Log Retention Method: overwrite-as-needed
Max Log Size: 16384
File Security:
/vol1/home
/vol1/dir1

Kerberos:
Max Clock Skew: 5
Max Ticket Age: 10
Max Renew Age: 7
Privilege Rights:
Take Ownership: usr1, usr2
Security Privilege: usr1, usr2
Change Notify: usr1, usr2
Registry Values:
Signing Required: false
Restrict Anonymous:
No enumeration of SAM accounts: true
No enumeration of SAM accounts and shares: false
Restrict anonymous access to shares and named pipes: true
Combined restriction for anonymous user: no-access
Restricted Groups:
gpr1
gpr2
Central Access Policy Settings:
Policies: cap1
cap2

```
[[IDc22668e3c71d7caeb5cf538994588727]]
= vserver cifs group-policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show group policy configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy show` command displays information about
group policy configuration for CIFS-enabled Vserver. It displays all or a
subset of the group policy configuration matching the criteria that you
specify.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays group policy
configuration only for the Vserver that you specify.

``[-status {enabled|disabled}]`` - Group Policy Status::
If you specify this parameter, the command displays group policy
configuration only for the Vservers that match the status you specify.

== Examples

The following example displays group policy configuration for all
Vservers:
```

cluster1::> vserver cifs group-policy show

Vserver GPO Status
---------- ----------
vs1 disabled

```
[[ID3094ee7ae8263567837e12d18512bfc1]]
= vserver cifs group-policy update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Apply group policy settings defined in Active Directory

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy update` command applies the group-policy
settings defined in Active Directory for the given Vserver. This command
is not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the CIFS-enabled Vserver to which the group-
policy settings be applied.

``[-force-reapply-all-settings {true|false}]`` - Force Re-apply All
Settings::
This parameter specifies whether to ignore all processing optimizations
and re-apply all settings. The default is false.

== Examples

The following example applies the group-policy settings defined in Active
Directory for Vserver vs1.
```

cluster1::> vserver cifs group-policy update -vserver vs1 -force-reapply-all-settings true

```
[[ID037720f96bb529db958fcff052e2f0d4]]
= vserver cifs group-policy central-access-policy show-applied
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show currently applied central access policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy central-access-policy show-applied` command
displays information about the central access policies assigned to
Vservers. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information about all CIFS servers:

* Vserver name
* Name of the central access policy
* SID
* Description
* Creation time
* Modification time
* Member rules



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only for
central access policies for the specified Vserver.

``[-name <TextNoCase>]`` - Name::
If you specify this parameter, the command displays information only for
central access policies that match the specified name.

``[-sid <windows sid>]`` - Identifier::
If you specify this parameter, the command displays information only for
```

```
central access policies that match the specified SID.

``[-description <text>]`` - Description::
If you specify this parameter, the command displays information only for
central access policies that match the specified description.

``[-ctime <Date>]`` - Creation Time::
If you specify this parameter, the command displays information only for
central access policies that match the specified creation time.

``[-mtime <Date>]`` - Modification Time::
If you specify this parameter, the command displays information only for
central access policies that match the specified modification time.

``[-rules <TextNoCase>,...]`` - Central Access Rules::
If you specify this parameter, the command displays information only for
central access policies that match the specified member rules.

== Examples

The following example displays information for all central access
policies:
```

cluster1::> vserver cifs group-policy central-access-policy show-applied

Vserver Name SID
--- --- ---
vs1 p1 S-1-17-3386172923-1132988875-3044489393-3993546205
Description: policy #1
Creation Time: Tue Oct 22 09:34:13 2013
Modification Time: Wed Oct 23 08:59:15 2013
Member Rules: r1

vs1 p2 S-1-17-1885229282-1100162114-134354072-822349040
Description: policy #2
Creation Time: Tue Oct 22 10:28:20 2013
Modification Time: Thu Oct 31 10:25:32 2013
Member Rules: r1
r2

2 entries were displayed.

```
[[ID90ae86614ae9b29ebda4f1db83cc2dcb]]
= vserver cifs group-policy central-access-policy show-defined
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Show applicable central access policies defined in the Active Directory

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy central-access-policy show-defined` command
displays information about the central access policies that are defined in
the Active Directory. The command output depends on the parameter or
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all CIFS
servers:

* Vserver name
* Name of the central access policy
* SID
* Description
* Creation time
* Modification time
* Member rules



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only for
central access policies for the specified Vserver.

``[-name <TextNoCase>]`` - Name::
If you specify this parameter, the command displays information only for
central access policies that match the specified name.

``[-sid <windows sid>]`` - Identifier::
If you specify this parameter, the command displays information only for
central access policies that match the specified SID.
```

```
``[-description <text>]`` - Description::
If you specify this parameter, the command displays information only for
central access policies that match the specified description.

``[-ctime <Date>]`` - Creation Time::
If you specify this parameter, the command displays information only for
central access policies that match the specified creation time.

``[-mtime <Date>]`` - Modification Time::
If you specify this parameter, the command displays information only for
central access policies that match the specified modification time.

``[-rules <TextNoCase>,...]`` - Central Access Rules::
If you specify this parameter, the command displays information only for
central access policies that match the specified member rules.

== Examples

The following example displays information for all central access
policies:
```

cluster1::> vserver cifs group-policy central-access-policy show-defined

Vserver Name SID
---------- -------------------- --
vs1 p1 S-1-17-3386172923-1132988875-3044489393-3993546205
Description: policy #1
Creation Time: Tue Oct 22 09:34:13 2013
Modification Time: Wed Oct 23 08:59:15 2013
Member Rules: r1

vs1 p2 S-1-17-1885229282-1100162114-134354072-822349040
Description: policy #2
Creation Time: Tue Oct 22 10:28:20 2013
Modification Time: Thu Oct 31 10:25:32 2013
Member Rules: r1
r2

2 entries were displayed.

```
[[ID467275c577772e7282280f1f54defe60]]
= vserver cifs group-policy central-access-rule show-applied
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Show currently applied central access rules

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy central-access-rule show-applied` command
displays information about the central access rules assigned to Vservers.
The command output depends on the parameter or parameters specified with
the command. If you do not specify any parameters, the command displays
the following information about all CIFS servers:

* Vserver name
* Name of the central access rule
* Description
* Creation time
* Modification time
* Current permissions
* Proposed permissions
* Target resources



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only for
central access rules for the specified Vserver.

``[-name <TextNoCase>]`` - Name::
If you specify this parameter, the command displays information only for
central access rules that match the specified name.

``[-description <text>]`` - Description::
If you specify this parameter, the command displays information only for
central access rules that match the specified description.
```

```
``[-ctime <Date>]`` - Creation Time::
If you specify this parameter, the command displays information only for
central access rules that match the specified creation time.

``[-mtime <Date>]`` - Modification Time::
If you specify this parameter, the command displays information only for
central access rules that match the specified modification time.

``[-effective <text>]`` - Effective Security Policy::
If you specify this parameter, the command displays information only for
central access rules that match the specified effective security policy.

``[-proposed <text>]`` - Proposed Security Policy::
If you specify this parameter, the command displays information only for
central access rules that match the specified proposed security policy.

``[-resource <text>]`` - Resource Condition::
If you specify this parameter, the command displays information only for
central access rules that match the specified resource condition.

== Examples

The following example displays information for all central access rules:
```

cluster1::> vserver cifs group-policy central-access-rule show-applied

Vserver Name
--------- ----------
vs1 r1
Description: rule #1
Creation Time: Tue Oct 22 09:33:48 2013
Modification Time: Tue Oct 22 09:33:48 2013
Current Permissions: O:SYG:SYD:AR(A;;FA;;;WD)
Proposed Permissions: O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)

vs1 r2
Description: rule #2
Creation Time: Tue Oct 22 10:27:57 2013
Modification Time: Tue Oct 22 10:27:57 2013
Current Permissions: O:SYG:SYD:AR(A;;FA;;;WD)
Proposed Permissions: O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)

2 entries were displayed.

```
[[ID59bb719da6e3a983a9a1fa1ecee375e8]]
= vserver cifs group-policy central-access-rule show-defined
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show applicable central access rules defined in the Active Directory

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy central-access-rule show-defined` command
displays information about the central access rules that are defined in
the Active Directory. The command output depends on the parameter or
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all CIFS
servers:

* Vserver name
* Name of the central access rule
* Description
* Creation time
* Modification time
* Current permissions
* Proposed permissions
* Target resources



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only for
central access rules for the specified Vserver.

``[-name <TextNoCase>]`` - Name::
If you specify this parameter, the command displays information only for
central access rules that match the specified name.
```

```
``[-description <text>]`` - Description::
If you specify this parameter, the command displays information only for
central access rules that match the specified description.

``[-ctime <Date>]`` - Creation Time::
If you specify this parameter, the command displays information only for
central access rules that match the specified creation time.

``[-mtime <Date>]`` - Modification Time::
If you specify this parameter, the command displays information only for
central access rules that match the specified modification time.

``[-effective <text>]`` - Effective Security Policy::
If you specify this parameter, the command displays information only for
central access rules that match the specified effective security policy.

``[-proposed <text>]`` - Proposed Security Policy::
If you specify this parameter, the command displays information only for
central access rules that match the specified proposed security policy.

``[-resource <text>]`` - Resource Condition::
If you specify this parameter, the command displays information only for
central access rules that match the specified resource condition.

== Examples

The following example displays information for all central access rules:
```

cluster1::> vserver cifs group-policy central-access-rule show-defined

Vserver Name
--- ---
vs1 r1
Description: rule #1
Creation Time: Tue Oct 22 09:33:48 2013
Modification Time: Tue Oct 22 09:33:48 2013
Current Permissions: O:SYG:SYD:AR(A;;FA;;;WD)
Proposed Permissions: O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)

vs1 r2
Description: rule #2
Creation Time: Tue Oct 22 10:27:57 2013
Modification Time: Tue Oct 22 10:27:57 2013
Current Permissions: O:SYG:SYD:AR(A;;FA;;;WD)
Proposed Permissions: O:SYG:SYD:(A;;FA;;;OW)(A;;FA;;;BA)(A;;FA;;;SY)

2 entries were displayed.

```
[[ID072c61bfa4179ed5d2b817afbe24290f]]
= vserver cifs group-policy restricted-group show-applied
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the applied restricted-group settings.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy restricted-group show-applied` command
displays settings of all the restricted groups applied to a Vserver.

If you do not specify any parameters, the command displays the following
information about all the restricted groups applied to all the Vservers in
the cluster.

* Group Policy Name: Specifies the name of the group policy.
* Version: Specifies the version of the group policy.
* Link: Specifies the level in which the group policy is configured.
Possible values are:

  * Local: Group policy is configured in ONTAP.
  * Site: Group policy is configured at the site level in the Domain
Controller.
  * Domain: Group policy is configured at the domain level in the Domain
Controller.
  * OrganizationalUnit: Group policy is configured at the OU level in the
Domain controller.
  * RSOP: Resultant set of policies derived from all the group policies
defined at various levels.


* Group Name: Specifies the name of a restricted group.
* Members: Specifies users and groups who belong to and who do not belong
to the restricted group.
* MemberOf: Specifies list of groups to which the restricted group is
added. A group can be a member of groups other than the groups listed
here.
```

```
== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays the restricted-group
information that has been applied to the specified Vserver.

``[-index <integer>]`` - Index::
If this parameter is specified, the command displays the specified index
for the group policy in the restricted group. The restricted-group
information for the group policy at the specified index.

``[-group-name <text>]`` - Group Name::
If this parameter is specified, the command displays the restricted-group
information for the specified group name.

``[-group-policy-name <text>]`` - Group Policy Name::
If this parameter is specified, the command displays the restricted-group
information for the specified group policy name.

``[-uuid <UUID>]`` - UUID::
If this parameter is specified, the command displays the restricted-group
information for the specified UUID of the group policy.

``[-version <integer>]`` - Version::
If this parameter is specified, the command displays the restricted-group
information for the specified version of the group policy.

``[-link <gpo-link>]`` - Link Type::
If this parameter is specified, the command displays the restricted-group
information for the specified link for the group policy.

``[-members <gpoUserGroup>,...]`` - Members, List of Users/groups::
If this parameter is specified, the command displays the restricted-group
information for the specified members of users and groups.

``[-member-of <gpoUserGroup>,...]`` - MemberOf, List of Groups::
If this parameter is specified, the command displays the restricted-group
information for the specified member of the group.
```

```
== Examples

The following example displays information about all restricted groups
that have been applied to a Vserver.
```

cluster1::> vserver cifs group-policy restricted-group show-applied

Vserver: vs_1

```
     Group Policy Name: gpo1
               Version: 16
                  Link: OrganizationalUnit
            Group Name: grp1
               Members: usr1
              MemberOf: GPO\g9
Group Policy Name: Resultant Set of Policy
               Version: 0
                  Link: RSOP
            Group Name: grp1
               Members: usr1
              MemberOf: GPO\g9
```

```
[[ID9b90c9d1b013602da813ce961800dc76]]
= vserver cifs group-policy restricted-group show-defined
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the defined restricted-group settings.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs group-policy restricted-group show-defined` command
displays settings of all the restricted groups defined in Domain
Controller for a Vserver.

If you do not specify any parameters, the command displays the following
information about all the restricted groups defined in Domain Controller
```

```
for all the Vservers in the cluster.

* Group Policy Name: Specifies the name of the group policy.
* Version: Specifies the version of the group policy.
* Link: Specifies the level in which the group policy is configured.
Possible values are:

  * Local: Group policy is configured in ONTAP.
  * Site: Group policy is configured at the site level in the Domain
Controller.
  * Domain: Group policy is configured at the domain level in the Domain
Controller.
  * OrganizationalUnit: Group policy is configured at the OU level in the
Domain Controller.
  * RSOP: Resultant set of policies derived from all the group policies
defined at various levels.


* Group Name: Specifies the name of a restricted group.
* Members: Specifies users and groups who belong to and who do not belong
to the restricted group.
* MemberOf: Specifies list of groups to which the restricted group is
added. A group can be a member of groups other than the groups listed
here.



== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays the restricted-group
information that is defined in Domain Controller for the specified
Vserver.

``[-index <integer>]`` - Index::
If this parameter is specified, the command displays the specified index
for the group policy in the restricted group. The restricted-group
information for the group policy at the specified index.
```

```
``[-group-name <text>]`` - Group Name::
If this parameter is specified, the command displays the restricted-group
information for the specified group name.

``[-group-policy-name <text>]`` - Group Policy Name::
If this parameter is specified, the command displays the restricted-group
information for the specified group policy name.

``[-uuid <UUID>]`` - UUID::
If this parameter is specified, the command displays the restricted-group
information for the specified UUID of the group policy.

``[-version <integer>]`` - Version::
If this parameter is specified, the command displays the restricted-group
information for the specified version of the group policy.

``[-link <gpo-link>]`` - Link Type::
If this parameter is specified, the command displays the restricted-group
information for the specified link for the group policy.

``[-members <gpoUserGroup>,...]`` - Members, List of Users/groups::
If this parameter is specified, the command displays the restricted-group
information for the specified members of users and groups.

``[-member-of <gpoUserGroup>,...]`` - MemberOf, List of Groups::
If this parameter is specified, the command displays the restricted-group
information for the specified member of the group.

== Examples

The following example displays information about all restricted groups
that are defined in Domain Controller for a Vserver.
```

cluster1::> vserver cifs group-policy restricted-group show-defined

Vserver: vs_1

```
     Group Policy Name: gpo1
               Version: 16
                  Link: OrganizationalUnit
            Group Name: grp1
               Members: usr1
              MemberOf: GPO\g9
Group Policy Name: Resultant Set of Policy
               Version: 0
                  Link: RSOP
            Group Name: grp1
               Members: usr1
              MemberOf: GPO\g9
```

```
[[ID762ccfb0b35f02c7124588ebf34b073d]]
= vserver cifs home-directory modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify attributes of CIFS home directories

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs home-directory modify` command modifies the CIFS home
directory configuration for a CIFS server. To use the home directory
option s (-is-home-dirs-access-for-admin-enabled or/and -is-home-dirs
-access-for-public-enabled ), a home directory share must be configured
with a dynamic share pattern preceded by a tilde(~). Valid dynamic share
patterns are ~%w and ~%d~%w. The pattern %u is not supported with the se
option s .

== Parameters

``-vserver <vserver>`` - Vserver::
This parameter specifies the name of the CIFS server for which you want to
modify the CIFS home directory configuration.

``[-is-home-dirs-access-for-admin-enabled {true|false}]`` - Is Home
Directory Access for Admin Enabled::
```

```
This optional parameter specifies whether a user with Windows
administrative privileges can connect to another user's home directory.
The default value for this parameter is ``_true_`` .

``[-is-home-dirs-access-for-public-enabled {true|false}]`` - Is Home
Directory Access for Public Enabled [.advaced-priv]#(privilege:
advanced)#::
This optional parameter specifies whether any user can connect to another
user's home directory. The default value for this parameter is ``_false_``
.

== Examples

The following example modifies the CIFS home directory configuration for
the Vserver "vs1". It enables users with Windows administrative privileges
to connect to another user's home directory , and enables any user to
connect to another user's home directory .
```

cluster1::> vserver cifs home-directory modify -vserver vs1 -is-home-dirs-access-for-admin-enabled true -is-home-dirs-access-for-public-enabled true

```
The following example shows the usage of the share creation pattern
~%d~%w.
```

cluster1::> vserver cifs share create -vserver vs1 -share-name %d%w -path %d/%w -share-properties homedirectory

```
The following example shows the usage of the share creation pattern ~%w.
```

cluster1::> vserver cifs share create -vserver vs1 -share-name ~%w -path %d/%w -share-properties homedirectory

```
[[IDf44d80ae2d8827b406d44afa3a983cdb]]
= vserver cifs home-directory show-user
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Home Directory Path for a User

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver cifs home-directory show-user` command prints the path of a
user's CIFS home directory. Use this command if multiple CIFS home
directory paths exist and you want to see which path holds the user's CIFS
home directory.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify this parameter, the command displays only the fields that
you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``-vserver <vserver>`` - Vserver::
Use this required parameter to specify the Vserver that contains the home
directory of the user specified with the required `-username` parameter.

``-username <text>`` - User Name::
Use this required parameter to locate the home directory of the specified
user. You must enter this parameter in the following format: user,
domain/user or cifs_server_name/user.

``[-path <text>]`` - Path::
If you specify this parameter, the command displays information about the
user's home directory with the specified path.

``[-share-name <text>]`` - Share Name::
If you specify this parameter, the command displays information about the
user's home directory with the specified home-directory share.

== Examples

The following command displays information about the home directory of
user gpo\rpuser1 belonging to Vserver vs1.
```

cluster1::> vserver cifs home-directory show-user -vserver vs1 -username gpo\rpuser1
Vserver : vs1
Username : GPO/rpuser1
ShareName Home Dir Path
--- ---
root /home/rpuser1
rpuser1 /home/rpuser1

GPOrpuser1 /home/GPO/rpuser1

```
The following command displays information about the home directory of
user gpo\rpuser1 belonging to Vserver vs1 at share path /home/rpuser1.
```

cluster1::> vserver cifs home-directory show-user -vserver vs1 -username gpo\rpuser1 -path /home/rpuser1
Vserver : vs1
Username : GPO/rpuser1
ShareName Home Directory Path
--- ---
root /home/rpuser1
rpuser1 /home/rpuser1
2 entries were displayed.

```
The following command displays information about the home directory of
user gpo\rpuser1 belonging to Vserver vs1 at share ~GPO~rpuser1.
```

cluster1::> vserver cifs home-directory show-user -vserver vs1 -username gpo\rpuser1 -share-name GPOrpuser1
Vserver : vs1
Username : GPO/rpuser1
ShareName Home Directory Path
--- ---
GPOrpuser1 /home/GPO/rpuser1

```
[[ID1ad31946bb472a6a99960ddc8164308b]]
= vserver cifs home-directory show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display home directory configurations

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs home-directory show` command displays the CIFS home
directory configuration for one or more Vservers.

== 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.

``[-vserver <vserver>]`` - Vserver::
If you specify this parameter, the command displays CIFS home directory
configuration for the specified Vserver.

``[-is-home-dirs-access-for-admin-enabled {true|false}]`` - Is Home
Directory Access for Admin Enabled::
If you specify this parameter, the command displays home directory
configuration for CIFS servers that have the specified setting.

``[-is-home-dirs-access-for-public-enabled {true|false}]`` - Is Home
Directory Access for Public Enabled [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays home directory
configuration for CIFS servers that have the specified setting.

== Examples

The following example lists the CIFS home directory configuration for a
Vserver on the cluster.
```

cluster1::> vserver cifs home-directory show -vserver vs1
Vserver: vs1
Is Home Directory Access for Admin Enabled: true

```
At the advanced privilege level or above, the output displays the
information below:
```

cluster1::*> vserver cifs options show
Vserver: vs1
Is Home Directory Access for Admin Enabled: true
Is Home Directory Access for Public Enabled: false

```
[[IDc3616a794c2e1ed6086958a30bf044bd]]
= vserver cifs home-directory search-path add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a home directory search path

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs home-directory search-path add` command adds a search
path to a CIFS home directory configuration.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver containing the CIFS home
directory configuration to which you want to add the search path.

``-path <text>`` - Path::
This parameter specifies the search path you want to add.

== Examples

The following example adds the path /home1 to the CIFS home directory
configuration on Vserver vs1.
```

cluster1::> vserver cifs home-directory search-path add -vserver vs1 -path /home1

```
[[IDa58aebc7b8bd8c885bb1bb94938a137b]]
= vserver cifs home-directory search-path remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a home directory search path

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs home-directory search-path remove` command removes a
search path from a CIFS home directory configuration.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver containing the CIFS home
directory configuration from which you want to remove the search path.

``-path <text>`` - Path::
This parameter specifies the search path you want to remove.

== Examples

The following example removes the path /home1 from the CIFS home directory
configuration on Vserver vs1.
```

cluster1::> vserver cifs home-directory search-path remove -vserver vs1 -path /home1

```
[[ID10327c04aba1d42c77e57933a2a151d4]]
= vserver cifs home-directory search-path reorder
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change the search path order used to find a match

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs home-directory search-path reorder` command moves a
search path to a new position in the search path order in the CIFS home
directory configuration for the CIFS-enabled Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS enabled Vserver for which you want to
reorder searches.

``-path <text>`` - Path::
This parameter specifies the search path you want to move.

``-to-position <integer>`` - Target Position::
This parameter specifies the new position of the search path in the search
path order.

== Examples

The following example moves the search path /home1 to position 1 in the
search path order for the CIFS home directory configuration on Vserver
vs1.
```

cluster1::> vserver cifs home-directory search-path reorder -vserver vs1 -path /home1 -to-position 1

```
[[ID95706ab030f9994590f62defdc4ca8b3]]
= vserver cifs home-directory search-path show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display home directory search paths

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs home-directory search-path show` command displays
information about the search paths that are in the home directory
configuration for the CIFS-enabled Vservers.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify this parameter, the command only displays the fields that
you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays home directory
configuration for the CIFS-enabled Vserver that you specify.

``[-path <text>]`` - Path::
If you specify this parameter, the command displays information only for
the search path that you specify.

== Examples

The following example displays information about search paths for all CIFS
home directories on all CIFS-enabled Vservers:
```

cluster1::> vserver cifs home-directory search-path show
Vserver Position Path
---------- -------- ----------
vs1 1 /home1

vs2 2 /home2

```
[[ID2cd8b35616864cf86f912efc93a0d63c]]
= vserver cifs options modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify CIFS options

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs options modify` command modifies CIFS options for a CIFS
server.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the CIFS server for which you want to
modify CIFS options.

``[-default-unix-user <text>]`` - Default UNIX User::
This optional parameter specifies the name of the default UNIX user for
the CIFS server.

``[-read-grants-exec {enabled|disabled}]`` - Read Grants Exec for Mode
Bits::
This optional parameter specifies whether the CIFS server does read grant
execution for mode bits.

``[-wins-servers <InetAddress>,...]`` - Windows Internet Name Service
(WINS) Addresses::
This optional parameter specifies a list of Windows Internet Name Server
(WINS) addresses for the CIFS server. You must specify the WINS servers
using an IP address. You can enter multiple WINS addresses as a comma-
delimited list.
+
NOTE: Use an IPv4 address because WINS over IPv6 is not supported.
+
```

```
``[-smb1-enabled {true|false}]`` - (DEPRECATED)-Enable SMB1 Protocol
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server negotiates the
SMB 1.0 version of the CIFS protocol. The default value for this parameter
is false.
+
NOTE: This parameter is deprecated because the SMB1 protocol is obsolete
and considered not secure. It might be removed in a future release.
+


``[-smb2-enabled {true|false}]`` - Enable all SMB2 Protocols [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server negotiates the
SMB 2 version of the CIFS protocol. The default value for this parameter
is true.

``[-smb3-enabled {true|false}]`` - Enable SMB3 Protocol [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server negotiates the
SMB 3 version of the CIFS protocol. The default value for this parameter
is true.

``[-smb31-enabled {true|false}]`` - Enable SMB3.1 Protocol [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server negotiates the
SMB 3.1 version of the CIFS protocol. The default value for this parameter
is true.

``[-max-mpx <integer>]`` - Maximum Simultaneous Operations per TCP
Connection [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of simultaneous
operations the CIFS server reports it can process per TCP connection. The
default value for this parameter is 255. This option only applies to SMB
1.0 protocol.

``[-shadowcopy-dir-depth <integer>]`` - Maximum Depth of Directories to
Shadow Copy [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum depth of directories on
which to create shadow copies in the CIFS server. The default for this
parameter is 5. The value 0 indicates that all sub-directories should be
shadow copied. This parameter is not supported for workgroup CIFS servers.
Directories and files within a FlexGroup will not be shadow copied because
FlexGroup volumes do not support shadow copy.

``[-copy-offload-enabled {true|false}]`` - Enable Copy Offload Feature
```

```
[.advaced-priv]#(privilege: advanced)#::
This optional parameter enables the Copy Offload feature in the CIFS
server. If set to false, the Copy Offload feature is disabled. The default
for this parameter is true.

``[-is-copy-offload-direct-copy-enabled {true|false}]`` - Is Direct-copy
Copy Offload Mechanism Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter enables the direct-copy mechanism for ODX copy
offload in the CIFS server. If set to false, the direct-copy mechanism is
disabled. The default for this parameter is true.  + The direct-copy
mechanism increases the performance of the copy offload operation when
Windows clients try to open the source file of a copy in a mode that
prevents the file being changed while the copy is in progress. If turned
off, regular copy offloading takes place.

``[-default-unix-group <text>]`` - Default UNIX Group::
This optional parameter specifies the name of the default UNIX group for
the CIFS server. If you do not specify a default UNIX group, the CIFS ACL
to NFSv4 ACL translation may result in incomplete NFSv4 ACL information.
This parameter is not supported by Vservers with FlexVol volumes.

``[-shadowcopy-enabled {true|false}]`` - Enable Shadow Copy Feature (VSS)
[.advaced-priv]#(privilege: advanced)#::
This optional parameter enables the Shadow Copy (VSS) feature in the CIFS
server when set to true. The VSS feature is disabled when set to false.
The default for this parameter is true. This parameter is not supported
for workgroup CIFS servers. Directories and files within a FlexGroup will
not be shadow copied because FlexGroup volumes do not support shadow copy.

``[-is-referral-enabled {true|false}]`` - Refer Clients to More Optimal
LIFs [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server automatically
refers clients to a data LIF local to the node which hosts the root of the
requested share. The default value for this parameter is false.

``[-is-local-auth-enabled {true|false}]`` - Is Local User Authentication
Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether local user authentication is
enabled for the CIFS server.

``[-is-local-users-and-groups-enabled {true|false}]`` - Is Local Users and
Groups Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the local users and groups
feature is enabled for the CIFS server.

``[-is-use-junctions-as-reparse-points-enabled {true|false}]`` - Is
```

```
Reparse Point Support Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server exposes junction
points to Windows clients as reparse points. The default value for this
parameter is true. This parameter is only active if the client has
negotiated use of the SMB 2 or SMB 3 protocol.

``[-is-exportpolicy-enabled {true|false}]`` - Is Export Policies for CIFS
Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server uses export
policies to control client access. The default value for this parameter is
false.

``[-is-unix-nt-acl-enabled {true|false}]`` - Is NT ACLs on UNIX Security-
style Volumes Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server has the NT ACLs
enabled on UNIX security-style volumes. The default value for this
parameter is true.

``[-is-trusted-domain-enum-search-enabled {true|false}]`` - Is Enumeration
of Trusted Domain and Search Capability Enabled [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports
enumeration of bidirectional trusted domains. It also supports the search
in all the bidirectional trusted domains when performing Windows user
lookups for UNIX user to Windows user name mapping. The default value is
true. This parameter is not supported for workgroup CIFS servers.

``[-client-session-timeout <integer>]`` - Idle Timeout Before CIFS Session
Disconnect (secs)::
This optional parameter specifies the amount of idle time (in seconds)
before a CIFS session is disconnected. The default value for this
parameter is 900 seconds.

``[-is-dac-enabled {true|false}]`` - Is Dynamic Access Control (DAC)
Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter enables the Dynamic Access Control (DAC) feature
in the CIFS server when set to true. The DAC feature is disabled when set
to false. The default for this parameter is false. This parameter is not
supported for workgroup CIFS servers.

``[-restrict-anonymous {no-restriction|no-enumeration|no-access}]`` -
Restrictions for Anonymous User [.advaced-priv]#(privilege: advanced)#::
This optional parameter controls the access restrictions of non-
authenticated sessions and applies the restrictions for the anonymous user
based on the permitted values. The default value for this parameter is no-
restriction. Permitted values for this option are:
```

```
+
  * no-restriction - This option specifies no access restriction for
anonymous users (default).
  * no-enumeration - This option specifies that only enumeration is
restricted.
  * no-access - This option specifies that access is restricted for
anonymous users.



``[-is-read-only-delete-enabled {enabled|disabled}]`` - Is Deletion of
Read-Only Files Enabled::
This optional parameter controls deletion of read-only files and
directories. NTFS delete semantics forbid deletion of a file or directory
when the read-only attribute is set. UNIX delete semantics ignore it,
focusing instead on parent directory permissions, which some applications
require. This option is used to select the desired behavior. By default
this option is disabled, yielding NTFS behavior.

``[-file-system-sector-size {512|4096 (in bytes)}]`` - Size of File System
Sector Reported to SMB Clients (bytes) [.advaced-priv]#(privilege:
advanced)#::
This optional parameter specifies the size of file system sector reported
to SMB clients (in bytes). The default value for this parameter is 4096.
Valid values are 512 and 4096.

``[-is-fake-open-enabled {true|false}]`` - Is Fake Open Support Enabled
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports fake
open requests. This parameter allows you to optimize the open and close
requests coming from SMB 2 clients. The default value for this parameter
is true.

``[-is-unix-extensions-enabled {true|false}]`` - Is UNIX Extensions
Enabled [.advaced-priv]#(privilege: advanced)#::
When set to true, this optional parameter enables the UNIX Extensions
feature in the CIFS server. If set to false, the UNIX Extensions feature
is disabled. The default for this parameter is false. UNIX Extensions
allows POSIX/UNIX style security to be displayed through the CIFS
protocol.

``[-is-search-short-names-enabled {true|false}]`` - Is Search Short Names
Support Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports
searching short names. A search query with this option enabled will try to
match 8.3 file names along with long file names. The default value for
```

```
this parameter is false.

``[-is-advanced-sparse-file-support-enabled {true|false}]`` - Is Advanced
Sparse File Support Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports the
advanced sparse file capabilities. This allows CIFS clients to query the
allocated ranges of a file and to write zeroes or free data blocks for
ranges of a file.

``[-is-fsctl-file-level-trim-enabled {true|false}]`` - Is Fsctl File Level
Trim Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether trim requests
(FSCTL_FILE_LEVEL_TRIM) are supported on the CIFS server.

``[-guest-unix-user <text>]`` - Map the Guest User to Valid UNIX User
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies that an unauthenticated user coming from
any untrusted domain can be mapped to a specified UNIX user for the CIFS
server. If the CIFS server cannot authenticate the user against a domain
controller for the home domain or a trusted domain or the local database,
and this option is enabled, the CIFS server considers the user as a guest
user and maps the user to the specified UNIX user. The UNIX user must be a
valid user.

``[-smb1-max-buffer-size <integer>]`` - Maximum Buffer Size Used for SMB1
Message [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum buffer size used for an SMB
1.0 message that the CIFS server can receive. If the LARGE_READ or
LARGE_WRITE capability is negotiated during session setup, then 'Read' or
'Write' SMB 1.0 operations are allowed to exceed the configured 'smb1-max-
buffer-size' value. This parameter does not have any effect on SMB 2 or
SMB 3 buffer size. The default value for this parameter is 65535. The
supported range for this parameter is 4356 through 65535.

``[-max-same-user-sessions-per-connection <integer>]`` - Maximum Same User
Sessions per TCP Connection [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of CIFS sessions that
can be set up by the same user per TCP connection. The default value of
this parameter is 2500. The maximum value of this parameter is 4294967295.

``[-max-same-tree-connect-per-session <integer>]`` - Maximum Same Tree
Connect per Session [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of CIFS tree connects
to the same share per CIFS session. The default value of this parameter is
5000. The maximum value of this parameter is 4294967295.
```

```
``[-max-opens-same-file-per-tree <integer>]`` - Maximum Opens on Same File
per Tree [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of existing opens on
the same file per CIFS tree. The default value of this parameter is 1000.
The maximum value of this parameter is 4294967295.

``[-max-watches-set-per-tree <integer>]`` - Maximum Watches Set per Tree
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of watches, also
known as change notifies, that can be set per CIFS tree. Tree here refers
to a share connect from a single client. The default value of this
parameter is 500. The maximum value of this parameter is 4294967295.

``[-is-admin-users-mapped-to-root-enabled {true|false}]`` - Map
Administrators to UNIX User 'root' [.advaced-priv]#(privilege:
advanced)#::
If this optional parameter is set to true, Windows users who are members
of the "BUILTIN\Administrators" group are mapped to UNIX user 'root'
unless a user who is a member of this group is explicitly mapped to a UNIX
user. If a Windows user is a member of the "BUILTIN\Administrators" group
and an explicit user mapping exists for that user, the explicit name
mapping takes precedence. If this parameter is set to false, users that
are members of the "BUILTIN\Administrators" group are not mapped to UNIX
'root'. The default value for this parameter is true.

``[-is-advertise-dfs-enabled {true|false}]`` - (DEPRECATED)-Enable DFS
Referral Advertisement [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to advertise DFS referral of the
CIFS protocol. The default value for this parameter is false. This option
is not applicable to SMB 1.0.
+
NOTE: This parameter is deprecated and may be removed in a future release
of ONTAP. The functionality provided by this parameter is now controlled
by the `-symlink-properties` parameter instead.
+


``[-is-path-component-cache-enabled {true|false}]`` - Is Path Component
Cache Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the path component cache is
enabled. The default value for this parameter is true.

``[-win-name-for-null-user <TextNoCase>]`` - Map Null User to Windows User
or Group [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies a valid Windows user or group name that
will be added to the CIFS credentials for a NULL user Session.
```

```
``[-is-hide-dotfiles-enabled {true|false}]`` - Is Hide Dot Files Enabled
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports hiding
dot files. Directory enumeration with this option enabled hides files and
directories that begin with a dot (".") . The default value for this
parameter is false.

``[-is-client-version-reporting-enabled {true|false}]`` - Is Client
Version Reporting Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, CIFS client version tracking information
is collected by AutoSupport. The default value of this parameter is true.

``[-is-client-dup-detection-enabled {true|false}]`` - Is Client Duplicate
Session Detection Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports
duplicate session detection. Duplicate sessions that come from the same
client with VcNumber of zero with this option enabled will be closed, and
is only applicable for SMB 1.0 clients. The default value for this
parameter is true.

``[-grant-unix-group-perms-to-others {true|false}]`` - Grant UNIX Group
Permissions to Others [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the incoming CIFS user who is
not the owner of the file, can be granted the group permission. If the
CIFS incoming user is not the owner of UNIX security-style file and this
option is set to true, then at all times the file's "group" permissions
are granted. If the CIFS incoming user is not the owner of UNIX security-
style file and this option is set to false, then the normal UNIX rules are
applicable to grant the permissions. The default value of this parameter
is false.

``[-is-multichannel-enabled {true|false}]`` - Is Multichannel Enabled
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports
Multichannel or not. The default value for this parameter is false.

``[-max-connections-per-session <integer>]`` - Maximum Connections Allowed
Per Multichannel Session [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of connections
allowed per Multichannel session. The default value for this parameter is
32.

``[-max-lifs-per-session <integer>]`` - Maximum LIFs Advertised Per
Multichannel Session [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of network interfaces
advertised per Multichannel session. The default value for this parameter
```

```
is 256.

``[-is-large-mtu-enabled {true|false}]`` - Is Large MTU Enabled [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports the SMB
2.1 "large MTU" feature. The default value for this parameter is true.

``[-is-netbios-over-tcp-enabled {true|false}]`` - Is NetBIOS over TCP
(port 139) Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports the
NetBIOS over TCP (port 139) feature. The default value for this parameter
is true.

``[-is-nbns-enabled {true|false}]`` - Is NBNS over UDP (port 137) Enabled
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the CIFS server supports the
NBNS protocol. The default value for this parameter is ``_false_`` .

``[-widelink-as-reparse-point-versions <CIFS Dialects>,...]`` - Protocol
Versions for Which Widelink Will Be Reported as Reparse Point [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the CIFS protocol versions for which the
widelink is reported as reparse point. The default value for this
parameter is ``_SMB1_`` .
+
NOTE: Any values entered for this parameter is replaced with the existing
values.
+


``[-max-credits <integer>]`` - Maximum Credits to Grant [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of outstanding
requests on a CIFS connection. The default value for this parameter is
128.

``[-is-inherit-modebits-with-nfsv4acl-enabled {true|false}]`` - Enable
Modebits on CIFS File Inheriting NFSv4 ACLs [.advaced-priv]#(privilege:
advanced)#::
This optional parameter specifies whether to set mode bits on the files
created by the cifs user that inherit NFSv4 acls. This parameter is not
supported for SMB1 clients.

``[-is-share-enum-permission-check-enabled {true|false}]`` - Check Share
Permission for NetShareEnumAll Request [.advaced-priv]#(privilege:
advanced)#::
```

```
If this parameter is set to ``_true_`` , the NetShareEnum call will only
respond with the shares the user has access to. The default value is
``_false_`` which means it will respond with all shares.

``[-is-backup-symlink-enabled {true|false}]`` - Preserve UNIX Symlink
During Backup through SMB [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the UNIX symlinks are preserved
during backup through SMB. The default value is false which means symlink
will not be preserved and it will be copied as regular file or directory.

== Examples

The following example modifies CIFS options for the Vserver "vs1". It
changes the default UNIX user, disables read grants exec, disables SMB2.x,
changes maximum multiplex count to 1124, changes the file system sector
size reported to SMB clients to 512, disables the direct-copy offload
mechanism for ODX copy offload, enables the UNIX Extensions feature,
disables fake open requests changes WINS servers to 192.168.11.112 and
changes the client session timeout to 6000.
```

cluster1::> vserver cifs options modify -vserver vs1
-default-unix-user pcuser -read-grants-exec disabled
-smb2-enabled false -max-mpx 1124 -file-system-sector-size
512 -is-copy-offload-direct-copy-enabled false
-is-unix-extensions-enabled true -is-fake-open-enabled false
-wins-servers 192.168.11.112 -client-session-timeout 6000

```
The following example modifies CIFS options for the Vserver "vs1". It
enables the advanced sparse file support .
```

cluster1::> vserver cifs options modify -vserver vs1
-is-advanced-sparse-file-support-enabled true

```
The following example modifies CIFS options for the Vserver "vs1". It
modifies limits for maximum opens on the same file, max sessions by the
same user, max tree connects per session, and max watches set.
```

cluster1::> vserver cifs options modify -vserver vs1
-max-same-user-sessions-per-connection 100
-max-same-tree-connect-per-session 100 -max-opens-same-file-per-tree 150
-max-watches-set-per-tree 200

```
The following example modifies CIFS options for the Vserver "vs1". It
modifies the option to disable the path component cache. .
```

cluster1::> vserver cifs options modify -vserver vs1
-is-path-component-cache-enabled false

```
The following example modifies CIFS options for the Vserver "vs1". It
modifies the option to disable CIFS client version tracking.
```

cluster1::> vserver cifs options modify -vserver vs1
-is-client-version-reporting-enabled false

```
The following example modifies CIFS option for the Vserver "vs1". It
modifies the option to enable granting of UNIX group permissions to
others.
```

cluster1::> vserver cifs options modify -vserver vs1
-grant-unix-group-perms-to-others true

```
[[IDf9311eebdbef02f07f6eb254cb7b044a]]
= vserver cifs options show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display CIFS options

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs options show` command displays the CIFS configuration
options for one or more Vservers.

== 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.
```

```
``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command only displays CIFS options for
the specified Vserver.

``[-default-unix-user <text>]`` - Default UNIX User::
If you specify this parameter, the command displays options for CIFS
server with the specified UNIX user.

``[-read-grants-exec {enabled|disabled}]`` - Read Grants Exec for Mode
Bits::
If this parameter is set to enabled, the command displays options for CIFS
servers that grant execution access when granting read access using mode
bits. If set to disabled, the command displays options for CIFS servers
that do not grant execution access when granting read access using mode
bits.

``[-wins-servers <InetAddress>,...]`` - Windows Internet Name Service
(WINS) Addresses::
If you specify this parameter, the command displays CIFS options only for
CIFS servers that use the specified Windows Internet Name Server (WINS)
addresses.

``[-smb1-enabled {true|false}]`` - (DEPRECATED)-Enable SMB1 Protocol
[.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where SMB 1.0 version of the CIFS protocol is negotiated. If set
to false, the command displays options for CIFS servers where SMB 1.0
version of the CIFS protocol is not negotiated.
+
NOTE: This parameter is deprecated because the SMB1 protocol is obsolete
and considered not secure. It might be removed in a future release.
+


``[-smb2-enabled {true|false}]`` - Enable all SMB2 Protocols [.advaced-
priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where SMB 2 version of the CIFS protocol is negotiated. If set to
false, the command displays options for CIFS servers where SMB 2 version
of the CIFS protocol is not negotiated.

``[-smb3-enabled {true|false}]`` - Enable SMB3 Protocol [.advaced-
priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where SMB 3 version of the CIFS protocol is negotiated. If set to
false, the command displays options for CIFS servers where SMB 3 version
```

```
of the CIFS protocol is not negotiated.

``[-smb31-enabled {true|false}]`` - Enable SMB3.1 Protocol [.advaced-
priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where SMB 3.1 version of the CIFS protocol is negotiated. If set
to false, the command displays options for CIFS servers where SMB 3.1
version of the CIFS protocol is not negotiated.

``[-max-mpx <integer>]`` - Maximum Simultaneous Operations per TCP
Connection [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options for CIFS
server with the specified maximum number of simultaneous operations the
CIFS server can process per TCP connection.

``[-shadowcopy-dir-depth <integer>]`` - Maximum Depth of Directories to
Shadow Copy [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
servers that are configured with the specified depth of directories on
which to create shadow copies.

``[-copy-offload-enabled {true|false}]`` - Enable Copy Offload Feature
[.advaced-priv]#(privilege: advanced)#::
If set to true, this command displays options only for CIFS servers where
the Copy Offload feature is enabled. If set to false, options are
displayed for CIFS servers where the Copy Offload feature is disabled.

``[-is-copy-offload-direct-copy-enabled {true|false}]`` - Is Direct-copy
Copy Offload Mechanism Enabled [.advaced-priv]#(privilege: advanced)#::
If set to true, this command displays options only for CIFS servers where
the direct-copy mechanism for ODX Copy Offload is enabled. If set to
false, options are displayed for CIFS servers where the direct- copy
offload mechanism is disabled.  + The direct-copy mechanism increases the
performance of the copy offload operation when Windows clients try to open
the source file of a copy in a mode that prevents the file being changed
while the copy is in progress. If turned off, regular copy offloading
takes place.

``[-default-unix-group <text>]`` - Default UNIX Group::
If you specify this parameter, the command displays options for CIFS
server with the specified default UNIX group.

``[-shadowcopy-enabled {true|false}]`` - Enable Shadow Copy Feature (VSS)
[.advaced-priv]#(privilege: advanced)#::
If set to true, this command displays options only for CIFS servers where
the Shadow Copy (VSS) feature is enabled. If set to false, options are
```

```
displayed for CIFS servers where the Shadow Copy (VSS) feature is
disabled.

``[-is-referral-enabled {true|false}]`` - Refer Clients to More Optimal
LIFs [.advaced-priv]#(privilege: advanced)#::
If set to true, the command displays options for the CIFS server where the
CIFS server automatically refers clients to a data LIF local to the node
which hosts the root of the requested share. If set to false, the command
displays options for the CIFS server where the mechanism, to automatically
refer the clients to data LIF local to the node which hosts the root of
the requested share, is disabled.

``[-is-local-auth-enabled {true|false}]`` - Is Local User Authentication
Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays CIFS options only
for CIFS servers where local user authentication is enabled. If set to
false, the command displays options for CIFS servers where local user
authentication is disabled.

``[-is-local-users-and-groups-enabled {true|false}]`` - Is Local Users and
Groups Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays CIFS options only
for CIFS servers where the local users and groups feature is enabled. If
set to false, the command displays options for CIFS servers where the
local users and groups feature is disabled.

``[-is-use-junctions-as-reparse-points-enabled {true|false}]`` - Is
Reparse Point Support Enabled [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command only displays CIFS options for
Vservers which have the specified reparse point setting.

``[-is-exportpolicy-enabled {true|false}]`` - Is Export Policies for CIFS
Enabled [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command only displays CIFS options for
Vservers which have the specified export policy setting.

``[-is-unix-nt-acl-enabled {true|false}]`` - Is NT ACLs on UNIX Security-
style Volumes Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command only displays CIFS options
for Vservers that have the NT ACLs on UNIX security-style volumes enabled.
If set to false, the command displays CIFS options for Vservers that have
the NT ACLS on UNIX security-style volumes disabled.

``[-is-trusted-domain-enum-search-enabled {true|false}]`` - Is Enumeration
of Trusted Domain and Search Capability Enabled [.advaced-
priv]#(privilege: advanced)#::
```

```
If this parameter is set to true, the command displays CIFS options only
for CIFS servers that support enumeration of bidirectional trusted domains
and that support searching in all bidirectional trusted domains when
performing Windows user lookups for UNIX user to Windows user name
mapping. If set to false, the command displays options for CIFS servers
that do not support enumeration of bidirectional trusted domains.

``[-client-session-timeout <integer>]`` - Idle Timeout Before CIFS Session
Disconnect (secs)::
If you specify this parameter, the command displays options only for CIFS
servers that are configured with the specified client session timeout
value (in seconds).

``[-is-dac-enabled {true|false}]`` - Is Dynamic Access Control (DAC)
Enabled [.advaced-priv]#(privilege: advanced)#::
If set to true, this command displays options only for CIFS servers where
the Dynamic Access Control (DAC) feature is enabled. If set to false,
options are displayed for CIFS servers where the Dynamic Access Control
(DAC) feature is disabled.

``[-restrict-anonymous {no-restriction|no-enumeration|no-access}]`` -
Restrictions for Anonymous User [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays CIFS options only for
CIFS servers that have the specified permitted value for the anonymous
user. Permitted values for this option are:
+
  * no-restriction - There is no access restriction for anonymous users.
  * no-enumeration - Only enumeration is restricted.
  * no-access - Access is restricted for anonymous users.



``[-is-read-only-delete-enabled {enabled|disabled}]`` - Is Deletion of
Read-Only Files Enabled::
If you specify this parameter, the command displays options only for CIFS
servers that have the specified is-read-only-delete-enabled setting.

``[-file-system-sector-size {512|4096 (in bytes)}]`` - Size of File System
Sector Reported to SMB Clients (bytes) [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays options only for CIFS
servers that are configured with the specified file system sector size (in
bytes).

``[-is-fake-open-enabled {true|false}]`` - Is Fake Open Support Enabled
[.advaced-priv]#(privilege: advanced)#::
```

```
If you set this parameter to true, the command displays options for CIFS
servers where fake open is enabled. If set to false, the command displays
options for CIFS servers where fake open is disabled.

``[-is-unix-extensions-enabled {true|false}]`` - Is UNIX Extensions
Enabled [.advaced-priv]#(privilege: advanced)#::
If set to true, this command displays options only for CIFS servers where
the UNIX Extensions feature is enabled. If set to false, options are
displayed for CIFS servers where the UNIX Extensions feature is disabled.
UNIX Extensions allows POSIX/UNIX style security to be displayed through
the CIFS protocol.

``[-is-search-short-names-enabled {true|false}]`` - Is Search Short Names
Support Enabled [.advaced-priv]#(privilege: advanced)#::
If you set this parameter to true, the command displays options for CIFS
servers where search short names is enabled. If set to false, the command
displays options for CIFS servers where search short names is disabled.

``[-is-advanced-sparse-file-support-enabled {true|false}]`` - Is Advanced
Sparse File Support Enabled [.advaced-priv]#(privilege: advanced)#::
If set to true, the command displays options for CIFS servers where the
advanced sparse file capabilities for CIFS are enabled. If set to false,
options are displayed for CIFS servers where the advanced sparse file
capabilities for CIFS are disabled.

``[-is-fsctl-file-level-trim-enabled {true|false}]`` - Is Fsctl File Level
Trim Enabled [.advaced-priv]#(privilege: advanced)#::
If set to true, the command displays options for all the CIFS servers
where trim requests (FSCTL_FILE_LEVEL_TRIM) are supported. If set to
false, options are displayed for all the CIFS servers where trim requests
(FSCTL_FILE_LEVEL_TRIM) are not supported.

``[-guest-unix-user <text>]`` - Map the Guest User to Valid UNIX User
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options for CIFS
server with the specified guest UNIX user.

``[-smb1-max-buffer-size <integer>]`` - Maximum Buffer Size Used for SMB1
Message [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
servers that are configured with the specified maximum buffer size value.

``[-max-same-user-sessions-per-connection <integer>]`` - Maximum Same User
Sessions per TCP Connection [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
server that are configured with the specified maximum same user session
```

```
per connection.

``[-max-same-tree-connect-per-session <integer>]`` - Maximum Same Tree
Connect per Session [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
server that are configured with the specified maximum same tree connects
per session.

``[-max-opens-same-file-per-tree <integer>]`` - Maximum Opens on Same File
per Tree [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
server that are configured with the specified maximum opens on same file
per tree.

``[-max-watches-set-per-tree <integer>]`` - Maximum Watches Set per Tree
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
server that are configured with the specified maximum watches set per
tree. Tree here refers to a share connect from a single client.

``[-is-admin-users-mapped-to-root-enabled {true|false}]`` - Map
Administrators to UNIX User 'root' [.advaced-priv]#(privilege:
advanced)#::
If you set this parameter to true, the command displays options for CIFS
servers where members of "BUILTIN\Adminstrators" group are mapped to UNIX
user 'root'. If set to false, the command displays options for CIFS
servers where members of the "BUILTIN\Adminstrators" group are not mapped
to UNIX user 'root'.

``[-is-advertise-dfs-enabled {true|false}]`` - (DEPRECATED)-Enable DFS
Referral Advertisement [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays CIFS options only
for CIFS servers where DFS referral advertisement is enabled. If set to
false, the command displays options for CIFS servers where DFS referral
advertisement is disabled. This option is not applicable to SMB 1.0.
+
NOTE: This parameter is deprecated and may be removed in a future release
of ONTAP. The functionality provided by this parameter is now controlled
by the `-symlink-properties` parameter instead.
+


``[-is-path-component-cache-enabled {true|false}]`` - Is Path Component
Cache Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where the path component cache is enabled. If set to false, the
```

```
command displays options for CIFS servers where the path component cache
is disabled.

``[-win-name-for-null-user <TextNoCase>]`` - Map Null User to Windows User
or Group [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
servers that are configured to add the specified windows user or group
into CIFS credentials for null sessions.

``[-is-hide-dotfiles-enabled {true|false}]`` - Is Hide Dot Files Enabled
[.advaced-priv]#(privilege: advanced)#::
When set to true, this optional parameter enables the Hide Dot Files
feature in the CIFS server. If set to false, the Hide Dot Files feature is
disabled. The default value for this parameter is false.

``[-is-client-version-reporting-enabled {true|false}]`` - Is Client
Version Reporting Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where CIFS client version tracking is enabled. If set to false,
the command displays options for CIFS servers where CIFS client version
tracking is disabled.

``[-is-client-dup-detection-enabled {true|false}]`` - Is Client Duplicate
Session Detection Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where client duplicate session detection is enabled. If set to
false, the command displays options for CIFS servers where client
duplicate session detection is not enabled.

``[-grant-unix-group-perms-to-others {true|false}]`` - Grant UNIX Group
Permissions to Others [.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays CIFS options only
for CIFS servers where grant unix group permissions to others feature is
enabled. If set to false, the command displays options for CIFS servers
where grant unix group permissions to others feature is disabled.

``[-is-multichannel-enabled {true|false}]`` - Is Multichannel Enabled
[.advaced-priv]#(privilege: advanced)#::
If this parameter is set to true, the command displays options for CIFS
servers where the multichannel is enabled. If set to false, the command
displays options for CIFS servers where the multichannel is disabled.

``[-max-connections-per-session <integer>]`` - Maximum Connections Allowed
Per Multichannel Session [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
server that are configured with the specified maximum connections allowed
```

```
per multichannel session.

``[-max-lifs-per-session <integer>]`` - Maximum LIFs Advertised Per
Multichannel Session [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
server that are configured with the specified maximum network interfaces
advertised per multichannel session.

``[-is-large-mtu-enabled {true|false}]`` - Is Large MTU Enabled [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
servers that are configured to support the SMB 2.1 "Large MTU" feature.

``[-is-netbios-over-tcp-enabled {true|false}]`` - Is NetBIOS over TCP
(port 139) Enabled [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
servers that are configured to support the NetBIOS over TCP (port 139)
feature.

``[-is-nbns-enabled {true|false}]`` - Is NBNS over UDP (port 137) Enabled
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays CIFS options only for
CIFS servers that use the specified setting for the NBNS protocol.

``[-widelink-as-reparse-point-versions <CIFS Dialects>,...]`` - Protocol
Versions for Which Widelink Will Be Reported as Reparse Point [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays CIFS options only for
the CIFS servers that matches the specified CIFS protocol versions for
which widelinks are reported as reparse points. If a list is entered,
entries are returned that matches all the specified versions.

``[-max-credits <integer>]`` - Maximum Credits to Grant [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays options only for CIFS
servers that are configured with the specified maximum credits.

``[-is-inherit-modebits-with-nfsv4acl-enabled {true|false}]`` - Enable
Modebits on CIFS File Inheriting NFSv4 ACLs [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, mode bits are set on the files created by
the cifs user that inherit nfsv4 acls.

``[-is-share-enum-permission-check-enabled {true|false}]`` - Check Share
Permission for NetShareEnumAll Request [.advaced-priv]#(privilege:
advanced)#::
```

```
If this parameter is set to true, the command only displays CIFS options
for Vservers that enabled share permission check for NetShareEnumAll
request. If set to false, options are displayed for CIFS servers that
disabled share permission check for NetShareEnumAll request.

``[-is-backup-symlink-enabled {true|false}]`` - Preserve UNIX Symlink
During Backup through SMB [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays options only for CIFS
servers that are configured to support preserving UNIX symlinks during
backup through SMB.

== Examples
```

```
The following example lists CIFS options for a Vserver on the cluster.
```

cluster1::> vserver cifs options show

Vserver: vs1
Client Session Timeout: 900
Default Unix Group: -
Default Unix User: pcuser
Guest Unix User: guestusers
Read Grants Exec: disabled
WINS Servers: -

```
At the advanced level, the output displays the information below.
```

cluster1::*> vserver cifs options show

Vserver: vs1
Client Session Timeout: 900
Copy Offload Enabled: true
Default Unix Group: -
Default Unix User: pcuser
Guest Unix User: -
Are Administrators mapped to 'root': true
Is Advanced Sparse File Support Enabled: true
Direct-Copy Copy Offload Enabled: true
Export Policies Enabled: false
Grant Unix Group Permissions to Others: true
Is Advertise DFS Enabled: true
Is Client Duplicate Session Detection Enabled: true
Is Client Version Reporting Enabled: true
Is DAC Enabled: false
Is Fake Open Support Enabled: true
Is Hide Dot Files Support Enabled: false
Is Large MTU Enabled: true
Is Local Auth Enabled: true
Is Local Users and Groups Enabled: true

Is Multichannel Enabled: false
Is NetBIOS over TCP (port 139) Enabled: true
Is Referral Enabled: false
Is Search Short Names Support Enabled: false
Is Trusted Domain Enumeration And Search Enabled: true
Is UNIX Extensions Enabled: false
Is Use Junction as Reparse Point Enabled: true
Max Multiplex Count: 255
Max Connections per Multichannel Session: 32
Max LIFs per Multichannel Session: 256
Max Same User Session Per Connection: 2500
Max Same Tree Connect Per Session: 5000
Max Opens Same File Per Tree: 1000
Max Watches Set Per Tree: 500
NBNS Interfaces : -
Is Path Component Cache Enabled: true
NT ACLs on UNIX Security Style Volumes Enabled: true
Read Grants Exec: disabled
Read Only Delete: disabled
Reported File System Sector Size: 4096
Restrict Anonymous: no-restriction
Shadowcopy Dir Depth: 5
Shadowcopy Enabled: true
SMB1 Enabled: true
Max Buffer Size for SMB1 Message: 65535
SMB2 Enabled: true
SMB3 Enabled: true
SMB3.1 Enabled: true
Map Null User to Windows User or Group: cifsGroup
WINS Servers: -
Report Widelink as Reparse Point Versions: SMB1

```
[[IDc1d0138c5c4a970b0a252ec33b24c3cb]]
= vserver cifs security modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify CIFS security settings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs security modify` command modifies CIFS server security
settings.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver whose CIFS security
settings you want to modify.

``[-kerberos-clock-skew <integer>]`` - Maximum Allowed Kerberos Clock
Skew::
This parameter specifies the maximum allowed Kerberos ticket clock skew in
minutes. The default setting is 5 minutes.

``[-kerberos-ticket-age <integer>]`` - Kerberos Ticket Lifetime::
This parameter specifies the Kerberos ticket lifetime in hours. The
default setting is 10 hours.

``[-kerberos-renew-age <integer>]`` - Maximum Kerberos Ticket Renewal
Days::
This parameter specifies the maximum Kerberos ticket renewal lifetime in
days. The default setting is 7 days.

``[-kerberos-kdc-timeout <integer>]`` - Timeout for Kerberos KDC
Connections (Secs)::
This parameter specifies the timeout for sockets on KDCs after which all
KDCs are marked as unreachable. The default setting is 3 seconds.

``[-is-signing-required {true|false}]`` - Require Signing for Incoming
CIFS Traffic::
This parameter specifies whether signing is required for incoming CIFS
traffic. The default setting is ``_false_`` .

``[-is-password-complexity-required {true|false}]`` - Require Password
Complexity for Local User Accounts::
This parameter specifies whether password complexity is required for CIFS
local users. If this parameter is set to ``_true_`` , password complexity
is required. If the value is set to ``_false_`` , password complexity is
not required. The default setting is ``_true_`` for CIFS servers.

``[-use-start-tls-for-ad-ldap {true|false}]`` - Use start_tls for AD LDAP
Connections::
This parameter specifies whether to use Start TLS over AD LDAP
connections. When enabled, the communication between the ONTAP LDAP Client
and the LDAP Server will be encrypted using Start TLS. Start TLS is a
mechanism to provide secure communication by using the TLS/SSL protocols.
The default setting is ``_false_`` .
+
NOTE:  Ensure right certificates are installed for CIFS home domain and
```

```
trusted domains.
+


``[-is-aes-encryption-enabled {true|false}]`` - (DEPRECATED)-Is AES-128
and AES-256 Encryption for Kerberos Enabled::
This parameter specifies whether to use Kerberos AES-128 and AES-256
encryption types for authentication. When enabled, and depending on
negotiation with the KDC service, it is possible for authentication
operations to use these encryption types. The default setting is
``_true_`` .
+
NOTE:  This parameter is deprecated and might be removed from a future
release.
+



``[-lm-compatibility-level {lm-ntlm-ntlmv2-krb|ntlm-ntlmv2-krb|ntlmv2-
krb|krb}]`` - LM Compatibility Level::
This parameter specifies the LM compatibility level. The default setting
is ``_lm-ntlm-ntlmv2-krb_`` (LM, NTLM, NTLMv2 and Kerberos).

``[-is-smb-encryption-required {true|false}]`` - Require SMB Encryption
for Incoming CIFS Traffic::
This parameter specifies whether SMB encryption is required when accessing
shares in the Vserver. When enabled and depending on negotiation during
session setup, it is possible that data transfers between the client and
the server are made secure by encrypting the SMB traffic. The default
setting is ``_false_`` .

``[-session-security-for-ad-ldap {none|sign|seal}]`` - Client Session
Security::
This parameter specifies the level of security to be used for LDAP
communications. The default setting is ``_none_`` .
+
LDAP Client Session Security can be one of the following:
+

+
  *  none - No Signing or Sealing.
  *  sign - Sign LDAP traffic.
  *  seal - Seal and Sign LDAP traffic.



``[-smb1-enabled-for-dc-connections {false|true|system-default}]`` -
(DEPRECATED)-SMB1 Enabled for DC Connections::
```

```
This parameter specifies whether SMB1 is enabled for use with connections
to domain controllers. The default setting is ``_system-default_`` .
+
SMB1 Enabled For DC Connections can be one of the following:
+

+
  *  false - SMB1 is not enabled.
  *  true - SMB1 is enabled.
  *  system-default - This sets the option to whatever is the default for
the release of ONTAP that is running. For this release it is: SMB1 is
enabled.


+
NOTE: This parameter is deprecated because the SMB1 protocol is obsolete
and considered not secure. It might be removed in a future release.
+


``[-smb2-enabled-for-dc-connections {false|true|system-default}]`` - SMB2
Enabled for DC Connections::
This parameter specifies whether SMB2 is enabled for use with connections
to domain controllers. The default setting is ``_system-default_`` .
+
SMB2 Enabled For DC Connections can be one of the following:
+

+
  *  false - SMB2 is not enabled.
  *  true - SMB2 is enabled.
  *  system-default - This sets the option to whatever is the default for
the release of ONTAP that is running. For this release it is: SMB2 is
enabled.



``[-referral-enabled-for-ad-ldap {true|false}]`` - LDAP Referral Chasing
Enabled For AD LDAP Connections::
This parameter specifies whether LDAP referral is enabled for AD LDAP
connections. The default setting is ``_false_`` .

``[-use-ldaps-for-ad-ldap {true|false}]`` - Use LDAPS for Secure Active
Directory LDAP Connections::
This parameter specifies whether to use LDAPS over AD LDAP connections.
When enabled, the communication between the ONTAP LDAP Client and the LDAP
Server will be encrypted using LDAPS and port 636 will be used. LDAPS is a
```

```
mechanism to provide secure communication by using the TLS/SSL protocols
and port 636. The default setting is ``_false_`` .
+
NOTE:  Ensure right certificates are installed for CIFS home domain and
trusted domains.
+


``[-encryption-required-for-dc-connections {true|false}]`` - Encryption is
required for DC Connection::
This parameter specifies whether encryption is required for use with
connections to domain controllers. The default setting is ``_false_`` .
+
Encryption required For DC Connections can be one of the following:
+

+
  *  false - Encryption is not required.
  *  true - Encryption is required.



``[-aes-enabled-for-netlogon-channel {true|false}]`` - AES session key
enabled for NetLogon channel::
This parameter specifies whether AES session key will be negotiated as
part of the NetLogon secure channel establishment. The default setting is
``_true_`` .

``[-try-channel-binding-for-ad-ldap {true|false}]`` - Try Channel Binding
For AD LDAP Connections::
This parameter specifies whether channel binding will be tried for AD LDAP
connections. The default setting is ``_true_`` . Channel binding will be
tried only if `-use-start-tls-for-ad-ldap` or `-use-ldaps-for-ad-ldap` is
enabled along with `-session-security-for-ad-ldap` set to either
``_sign_`` or ``_seal_`` .

``[-advertised-enc-types <CIFS Kerberos Encryption Type>,...]`` -
Encryption Types Advertised to Kerberos::
Encryption types advertised to Kerberos.The default setting is ` aes-256`
,` aes-128` ,` rc4` ,`des` .

== Examples

The following example makes the following changes: the Kerberos clock skew
is set to 3 minutes, the Kerberos ticket lifetime to 8 hours and it makes
signing required for Vserver "vs1".
```

cluster1::> vserver cifs security modify -vserver vs1 -kerberos-clock-skew 3 -kerberos-ticket-age 8 -is-signing -required true
cluster1::> vserver cifs security show
Vserver: vs1
Kerberos Clock Skew: 3 minutes
Kerberos Ticket Age: 8 hours
Kerberos Renewal Age: 7 days
Kerberos KDC Timeout: 3 seconds
Is Signing Required: true
Is Password Complexity Required: true
Use start_tls For AD LDAP connection: false
Is AES Encryption Enabled: true
LM Compatibility Level: krb
Is SMB Encryption Required: false
Client Session Security: none
SMB1 Enabled For DC Connections: system-default
SMB2 Enabled For DC Connections: system-default
LDAP Referral Chasing Enabled For AD LDAP Connections: false
Use LDAPS for AD LDAP Connections: true
Encryption required For DC Connections: false
AES enabled for Netlogon channel: false
Try Channel Binding For AD LDAP Connections: true
Encryption Types Advertised to Kerberos: aes-256, aes-128, des, rc4

```
[[IDc29df75e3bf1babca8899242b4d4f825]]
= vserver cifs security show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display CIFS security settings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs security show` command displays information about CIFS
server security settings.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields parameter, the command only displays the fields
that you specify.

| ``[-instance  ]`` }::
```

```
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver::
This parameter specifies the name of the Vserver whose CIFS security
settings you want to display.

``[-kerberos-clock-skew <integer>]`` - Maximum Allowed Kerberos Clock
Skew::
If this parameter is specified, the command displays information only
about the security settings that match the specified Kerberos ticket clock
skew.

``[-kerberos-ticket-age <integer>]`` - Kerberos Ticket Lifetime::
If this parameter is specified, the command displays information only
about the security settings that match the specified Kerberos ticket age.

``[-kerberos-renew-age <integer>]`` - Maximum Kerberos Ticket Renewal
Days::
If this parameter is specified, the command displays information only
about the security settings that match the specified Kerberos renewal age.

``[-kerberos-kdc-timeout <integer>]`` - Timeout for Kerberos KDC
Connections (Secs)::
If this parameter is specified, the command displays information only
about the security settings that match the specified Kerberos KDC timeout.

``[-realm <text>]`` - Kerberos Realm::
If this parameter is specified, the command displays information only
about the security settings that match the specified Kerberos realm.

``[-kdc-ip <text>,...]`` - KDC IP Address::
If this parameter is specified, the command displays information only
about the security settings that match the specified KDC IP address.

``[-kdc-name <text>,...]`` - KDC Name::
If this parameter is specified, the command displays information only
about the security settings that match the specified KDC name.

``[-site <text>,...]`` - KDC Site::
If this parameter is specified, the command displays information only
about the security settings that match the specified Windows site.

``[-is-signing-required {true|false}]`` - Require Signing for Incoming
CIFS Traffic::
This parameter specifies whether signing is required for incoming CIFS
traffic. If this parameter is specified, the command displays information
```

```
only about the security settings that match the specified value for is-
signing-required.

``[-is-password-complexity-required {true|false}]`` - Require Password
Complexity for Local User Accounts::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where password
complexity for local user accounts is required. If set to ``_false_`` ,
the command displays security configuration information for CIFS servers
where password complexity for local user accounts is not required.

``[-use-start-tls-for-ad-ldap {true|false}]`` - Use start_tls for AD LDAP
Connections::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where Start TLS
is used for communication with the AD LDAP Server. If set to ``_false_`` ,
the command displays CIFS security configuration information only for CIFS
servers where Start TLS is not used for communication with the AD LDAP
Server.

``[-is-aes-encryption-enabled {true|false}]`` - (DEPRECATED)-Is AES-128
and AES-256 Encryption for Kerberos Enabled::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where AES-128 and
AES-256 encryption types for Kerberos are enabled. If set to ``_false_`` ,
the command displays security configuration information for CIFS servers
where AES-128 and AES-256 encryption types for Kerberos are disabled.
+
NOTE: This parameter is deprecated and may be removed from a future
release.
+



``[-lm-compatibility-level {lm-ntlm-ntlmv2-krb|ntlm-ntlmv2-krb|ntlmv2-
krb|krb}]`` - LM Compatibility Level::
If this parameter is specified, the command displays information only
about the security settings that match the specified LM compatibility
level.

``[-is-smb-encryption-required {true|false}]`` - Require SMB Encryption
for Incoming CIFS Traffic::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where SMB
encryption is required. If set to ``_false_`` , the command displays
security configuration information for CIFS servers where SMB encryption
is not required.
```

```
``[-session-security-for-ad-ldap {none|sign|seal}]`` - Client Session
Security::
If this parameter is set to ``_seal_`` , the command displays CIFS
security configuration information only for CIFS servers where both
signing and sealing are required for LDAP communications. If set to
``_sign_`` , the command displays security configuration information for
CIFS servers where only signing is required for LDAP communications. If
set to ``_none_`` , the command displays security configuration
information for CIFS servers where no security is required for LDAP
communications.

``[-smb1-enabled-for-dc-connections {false|true|system-default}]`` -
(DEPRECATED)-SMB1 Enabled for DC Connections::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where SMB1 is
enabled for use with connections to domain controllers. If set to
``_false_`` , the command displays security configuration information for
CIFS servers where SMB1 is not enabled for use with connections to domain
controllers. If set to ``_system-default_`` , the command displays
security configuration information for CIFS servers where the system-
default setting (SMB1 enabled) is used for connections to domain
controllers.
+
NOTE: This parameter is deprecated because the SMB1 protocol is obsolete
and considered not secure. It might be removed in a future release.
+


``[-smb2-enabled-for-dc-connections {false|true|system-default}]`` - SMB2
Enabled for DC Connections::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where SMB2 is
enabled for use with connections to domain controllers. If set to
``_false_`` , the command displays security configuration information for
CIFS servers where SMB2 is not enabled for use with connections to domain
controllers. If set to ``_system-default_`` , the command displays
security configuration information for CIFS servers where the system-
default setting (SMB2 enabled) is used for connections to domain
controllers.

``[-referral-enabled-for-ad-ldap {true|false}]`` - LDAP Referral Chasing
Enabled For AD LDAP Connections::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where LDAP
referral is enabled for AD LDAP connections. If set to ``_false_`` , the
command displays security configuration information for CIFS servers where
```

```
LDAP referral is not enabled for AD LDAP connections.

``[-use-ldaps-for-ad-ldap {true|false}]`` - Use LDAPS for Secure Active
Directory LDAP Connections::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where LDAPS is
used for communication with the AD LDAP Server. If set to ``_false_`` ,
the command displays CIFS security configuration information only for CIFS
servers where LDAPS is not used for communication with the AD LDAP Server.

``[-encryption-required-for-dc-connections {true|false}]`` - Encryption is
required for DC Connection::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where encryption
is required for use with connections to domain controllers. If set to
``_false_`` , the command displays security configuration information for
CIFS servers where encryption is not required for use with connections to
domain controllers.

``[-aes-enabled-for-netlogon-channel {true|false}]`` - AES session key
enabled for NetLogon channel::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where AES session
key is used for Netlogon secure channel. If set to ``_false_`` , the
command displays CIFS security configuration information only for CIFS
servers where AES session key is not used for Netlogon secure channel.

``[-try-channel-binding-for-ad-ldap {true|false}]`` - Try Channel Binding
For AD LDAP Connections::
If this parameter is set to ``_true_`` , the command displays CIFS
security configuration information only for CIFS servers where channel
binding is tried for AD LDAP connections. If set to ``_false_`` , the
command displays CIFS security configuration information only for CIFS
servers where channel binding is not tried for AD LDAP connections.

``[-advertised-enc-types <CIFS Kerberos Encryption Type>,...]`` -
Encryption Types Advertised to Kerberos::
If this parameter is specified, the command displays information only
about the security settings that match the specified advertised encryption
types.

== Examples

The following example displays CIFS server security settings.
```

cluster1::> vserver cifs security show

Vserver: vs1
Kerberos Clock Skew: 3 minutes
Kerberos Ticket Age: 8 hours
Kerberos Renewal Age: 7 days
Kerberos KDC Timeout: 3 seconds
Is Signing Required: true
Is Password Complexity Required: true
Use start_tls For AD LDAP connection: false
Is AES Encryption Enabled: false
LM Compatibility Level: krb
Is SMB Encryption Required: false
Client Session Security: none
SMB1 Enabled For DC Connections: system-default
SMB2 Enabled For DC Connections: system-default
LDAP Referral Chasing Enabled For AD LDAP Connections: false
Use LDAPS for AD LDAP Connections: true
Encryption required For DC Connections: false
AES session key enabled for NetLogon channel: false
Try Channel Binding For AD LDAP Connections: true
Encryption Types Advertised to Kerberos: des, rc4

```
The following example displays the Kerberos clock skew for all Vservers.
```

cluster1::> vserver cifs security show -fields kerberos-clock-skew
vserver kerberos-clock-skew
------- ------------------
vs1 5

```
[[ID4d93d5226e70c6edeba3d9f5a4ff7278]]
= vserver cifs session close
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Close an open CIFS session

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs session close` command closes the specified CIFS
sessions.

== Parameters
```

```
``-node {<nodename>|local}`` - Node::
If you specify this parameter, the command will close all the opened CIFS
sessions on the specified node.

``-vserver <vserver name>`` - Vserver::
If you specify this parameter, the command will close all the opened CIFS
sessions on the specified CIFS-enabled Vserver.

``-session-id <integer>`` - Session ID::
If you specify this parameter, the command will close the open CIFS
session that matches the specified session ID.

``[-connection-id <integer>]`` - Connection ID::
If you specify this parameter, the command will close all the opened CIFS
sessions that match the specified connection ID.

``[-lif-address <IP Address>]`` - Incoming Data LIF IP Address::
If you specify this parameter, the command will close all the opened CIFS
sessions that are established through the specified data LIF IP address.

``[-address <IP Address>]`` - Workstation IP address::
If you specify this parameter, the command will close all the opened CIFS
sessions that are opened from the specified IP address.

``[-auth-mechanism <Authentication Mechanism>]`` - Authentication
Mechanism::
If you specify this parameter, the command will close all the opened CIFS
sessions that used the specified authentication mechanism. The
authentication mechanism can include one of the following:
+
  *  NTLMv1 - NTLMv1 authentication mechanism
  *  NTLMv2 - NTLMv2 authentication mechanism
  *  Kerberos - Kerberos authentication mechanism
  *  Anonymous - Anonymous authentication mechanism



``[-windows-user <TextNoCase>]`` - Windows User::
If you specify this parameter, the command will close all the opened CIFS
sessions that are established for the specified CIFS user. The acceptable
format for CIFS user is [domain]\user.

``[-unix-user <text>]`` - UNIX User::
If you specify this parameter, the command will close all the opened CIFS
sessions that are established for the specified UNIX user.

``[-protocol-version <CIFS Dialects>]`` - Protocol Version::
```

```
If you specify this parameter, the command will close all the opened CIFS
sessions that are established over the specified version of CIFS protocol.
The protocol version can include one of the following:
+
  * SMB1 - SMB 1.0
  * SMB2 - SMB 2.0
  * SMB2_1 - SMB 2.1
  * SMB3 - SMB 3.0
  * SMB3_1 - SMB 3.1



``[-continuously-available <CIFS Open File Protection>]`` - Continuously
Available::
If you specify this parameter, the command will close all the opened CIFS
sessions with open files that have the specified level of continuously
available protection. The open files are "continuously available" if they
are opened from an SMB 3 client through a share with the
"continuously_available" property set. These open files are capable of
non-disruptively recovering from takeover and giveback as well as general
aggregate relocation between partners in a high-availability relationship.
This is in addition to the traditional SMB 2 capability allowing clients
to recover from LIF migration and other brief network interruptions.
+
NOTE: The CA protection levels depict the continuous availability at the
connection level so it might not be accurate for a session if the
connection has multiple sessions. Streams opened through a continuously
available share are permitted, but are not currently made continuously
available. Directories may be opened through a continuously available
share, but, by design, will not appear continuously available as clients
do not open them that way. These protection levels are applicable to the
sessions on read/write volumes residing on storage failover aggregates.
+

+
The continuously available status can be one of the following:
+

+
  * No - The session contains one or more open file but none of them are
continuously available.
  * Yes - The session contains one or more open files and all of them are
continuously available.
  * Partial - The session contains at least one continuously available
open file but other open files that are not.
```

```
``[-is-session-signed {true|false}]`` - Is Session Signed::
If you specify this parameter, the command will close all the opened CIFS
sessions that are established with the specified SMB signing option.

``[-smb-encryption-status {unencrypted|encrypted|partially-encrypted}]`` -
SMB Encryption Status::
If you specify this parameter, the command will close all the opened CIFS
sessions that are established over the specified SMB encryption status.
+
The SMB encryption status can be one of the following:
+

+
  * unencrypted - The CIFS session is not encrypted.
  * encrypted - The CIFS session is fully encrypted. Vserver level
encryption is enabled and encryption happens for the entire session.
  * partially-encrypted - The CIFS session is partially encrypted. Share
level encryption is enabled and encryption is initiated when the tree-
connect occurs.



== Examples

The following example closes all open CIFS sessions on all the nodes with
protocol-version SMB2:
```

cluster1::> cifs session close -node * -protocol-version SMB2
2 entries were acted on.

```
The following example closes all open CIFS sessions for all Vservers on
node node1:
```

cluster1::> cifs session close -node node1 -vserver *
3 entries were acted on.

```
[[ID1b588c24c073bb91fd7c32a9e3f0af0b]]
= vserver cifs session show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Display established CIFS sessions

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs session show` command displays information about
established CIFS sessions. The command output depends on the parameter or
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all CIFS
sessions:

* Node name
* Vserver name
* CIFS connection ID
* CIFS session ID
* Workstation IP address
* CIFS user name
* CIFS open files
* Session idle time

 You can specify additional parameters to display only information that
matches those parameters. For example, to display information only about
CIFS sessions established on connection ID 2012, run the command with the
`-connection-id parameter set to ` 2012.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify this parameter, the command only displays the fields that
you specify.

| ``[-active-volumes  ]``::
If you specify this parameter, the command displays the list of Volumes
that session has been connected.

| ``[-show-win-unix-creds  ]``::
If you specify this parameter along with a valid session-id, the command
displays Windows and UNIX credentials along with the detailed information
about matching CIFS sessions.

| ``[-instance  ]`` }::
If you specify this parameter, the command displays detailed information
about matching CIFS sessions.

``[-node {<nodename>|local}]`` - Node::
```

```
If you specify this parameter, the command displays information about the
CIFS sessions on the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information about CIFS
sessions on the specified CIFS-enabled Vserver.

``[-session-id <integer>]`` - Session ID::
If you specify this parameter, the command displays information about the
CIFS session that match the specified session ID.

``[-connection-id <integer>]`` - Connection ID::
If you specify this parameter, the command displays information about CIFS
sessions that match the specified connection ID.

``[-lif-address <IP Address>]`` - Incoming Data LIF IP Address::
If you specify this parameter, the command displays information about CIFS
sessions that are established through the specified data LIF IP address.

``[-address <IP Address>]`` - Workstation IP address::
If you specify this parameter, the command displays information about CIFS
sessions that are opened from the specified IP address.

``[-auth-mechanism <Authentication Mechanism>]`` - Authentication
Mechanism::
If you specify this parameter, the command displays information about CIFS
sessions that used the specified authentication mechanism. The
authentication mechanism can include one of the following:
+
  *  None - Could not authenticate
  *  NTLMv1 - NTLMv1 authentication mechanism
  *  NTLMv2 - NTLMv2 authentication mechanism
  *  Kerberos - Kerberos authentication mechanism
  *  Anonymous - Anonymous authentication mechanism



``[-windows-user <TextNoCase>]`` - Windows User::
If you specify this parameter, the command displays information about CIFS
sessions that are established for the specified CIFS user. The acceptable
format for CIFS user is [domain]\user.

``[-unix-user <text>]`` - UNIX User::
If you specify this parameter, the command displays information about CIFS
sessions that are established for the specified UNIX user.

``[-shares <integer>]`` - Open Shares::
```

```
If you specify this parameter, the command displays information about CIFS
sessions that have the specified number of CIFS shares opened.

``[-files <integer>]`` - Open Files::
If you specify this parameter, the command displays information about CIFS
sessions that have the specified number of regular CIFS files opened.

``[-other <integer>]`` - Open Other::
If you specify this parameter, the command displays information about CIFS
sessions that have the specified number of special CIFS files opened such
as streams or directories.

``[-connected-time <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Connected Time::
If you specify this parameter, the command displays information about CIFS
sessions that are established for the specified time duration.

``[-idle-time <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Idle
Time::
If you specify this parameter, the command displays information about CIFS
sessions on which there is no activity for the specified time duration.

``[-protocol-version <CIFS Dialects>]`` - Protocol Version::
If you specify this parameter, the command displays information about CIFS
sessions that are established over the specified version of CIFS protocol.
The protocol version can include one of the following:
+
  * SMB1 - SMB 1.0
  * SMB2 - SMB 2.0
  * SMB2_1 - SMB 2.1
  * SMB3 - SMB 3.0
  * SMB3_1 - SMB 3.1



``[-continuously-available <CIFS Open File Protection>]`` - Continuously
Available::
If you specify this parameter, the command displays information about CIFS
sessions with open files that have the specified level of continuously
available protection. The open files are "continuously available" if they
are opened from an SMB 3 client through a share with the
"continuously_available" property set. These open files are capable of
non-disruptively recovering from takeover and giveback as well as general
aggregate relocation between partners in a high-availability relationship.
This is in addition to the traditional SMB 2 capability allowing clients
to recover from LIF migration and other brief network interruptions.
```

```
+
NOTE: The CA protection levels depict the continuous availability at the
connection level so it might not be accurate for a session if the
connection has multiple sessions. Streams opened through a continuously
available share are permitted, but are not currently made continuously
available. Directories may be opened through a continuously available
share, but, by design, will not appear continuously available as clients
do not open them that way. These protection levels are applicable to the
sessions on read/write volumes residing on storage failover aggregates.
+

+
The continuously available status can be one of the following:
+

+
  * No - The session contains one or more open file but none of them are
continuously available.
  * Yes - The session contains one or more open files and all of them are
continuously available.
  * Partial - The session contains at least one continuously available
open file but other open files that are not.



``[-is-session-signed {true|false}]`` - Is Session Signed::
If you specify this parameter, the command displays information about CIFS
sessions that are established with the specified SMB signing option.

``[-user-type {local-user|domain-user|guest-user|anonymous-user|hybrid-
user}]`` - User Authenticated as::
If you specify this parameter, the command displays information about CIFS
sessions that are established for the specified user type. The user type
can include one of the following:
+
  *  local-user - Authenticated as a local CIFS user
  *  domain-user - Authenticated as a domain user
  *  guest-user - Authenticated as a guest user
  *  anonymous-user - Authenticated as an anonymous or null user



``[-netbios-name <text>]`` - NetBIOS Name::
If you specify this parameter, the command displays information about CIFS
sessions that are established with the specified NetBIOS Name.

``[-smb-encryption-status {unencrypted|encrypted|partially-encrypted}]`` -
```

```
SMB Encryption Status::
If you specify this parameter, the command displays information about CIFS
sessions that are established with the specified SMB encryption status.
+
The SMB encryption status can be one of the following:
+

+
  * unencrypted - The CIFS session is not encrypted.
  * encrypted - The CIFS session is fully encrypted. Vserver level
encryption is enabled and encryption happens for the entire session.
  * partially-encrypted - The CIFS session is partially encrypted. Share
level encryption is enabled and encryption is initiated when the tree-
connect occurs.



``[-connection-count <integer>]`` - Connection Count::
If you specify this parameter, the command displays information about CIFS
sessions that have the specified number of CIFS connections.

``[-is-large-mtu-enabled {true|false}]`` - Is Large MTU Enabled::
If you specify this parameter, the command displays information about CIFS
sessions that are established with the specified Large MTU option.

``[-vol-names <volume name>,...]`` - Volumes List::
If you specify this parameter, the command displays information about CIFS
sessions that are established with the specified volume names.

``[-share-names <Share>,...]`` - Open Shares Lists::
If you specify this parameter, the command displays information about CIFS
sessions that are established including the specified share names.

== Examples

The following example displays information about all CIFS sessions:
```

cluster1::> vserver cifs session show
Node: node1
Vserver: vs1
Connection Session Open Idle Connection
ID ID Workstation Windows User Files Time Count
---------- ------- --------------- --------------- --------- -------------- -----------------
127834 1 172.17.193.172 CIFSQA\ 2 22s 4
Administrator

```
The following example displays information about a CIFS session with
session-id 1.
```

cluster1::> vserver cifs session show -session-id 1 -instance
Node: node1
Vserver: vs1
Session ID: 1
Connection ID: 127834
Incoming Data LIF IP Address: 10.53.13.224
Workstation: 172.17.193.172
Authentication Mechanism: NTLMv2
Windows User: CIFSQA\Administrator
UNIX User: root
Open Shares: 2
Open Files: 2
Open Other: 0
Connected Time: 2d 17h 58m 5s
Idle Time: 22s
Protocol Version: SMB3
Continuously Available: No
Is Session Signed: true
User Authenticated as: domain-user
NetBIOS Name: ALIAS1
SMB Encryption Status: encrypted
Connection Count: 4
Windows Unix Credentials: -
Active Volumes: vol1,fg

```
[[IDff50eb96abde78bf9f7dd883f1dd30d8]]
= vserver cifs session file close
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Close an open CIFS file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs session file close` command closes the specified open
CIFS file.

== Parameters

``-node {<nodename>|local}`` - Node::
If you specify this parameter, the command will close all the opened CIFS
files on the specified node.

``-vserver <vserver name>`` - Vserver::
If you specify this parameter, the command will close all the opened CIFS
files on the specified CIFS-enabled Vserver.

``-file-id <integer>`` - File ID::
If you specify this parameter, the command will close the opened CIFS file
that matches the specified file ID.

``[-connection-id <integer>]`` - Connection ID::
If you specify this parameter, the command will close all the opened CIFS
files connected on the specified connection ID.

``[-session-id <integer>]`` - Session ID::
If you specify this parameter, the command will close all the opened CIFS
files connected on the specified session ID.

== Examples

The following example closes all the opened CIFS files that are connected
to the data LIFs of Vserver vs1 on the node node1 with the connection-id
1:
```

cluster1::> vserver cifs session file close -node node1 -vserver vs1 -connection-id 1
5 entries were acted on.

```
The following example closes all the opened CIFS files on the node node1
with the file-id 1, connection-id 1 and the session-id 1:
```

cluster1::> vserver cifs session file close -node node1 -file-id 1 -connection-id 1 -session-id 1
2 entries were acted on.

```
[[ID8f290e128d3175e360755ef1824a89ca]]
= vserver cifs session file show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display opened CIFS files

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs session file show` command displays information about
all open CIFS files. The command output depends on the parameter or
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all open
CIFS files:

* Node name
* Vserver name
* CIFS connection ID
* CIFS session ID
* CIFS file ID
* CIFS file type
* CIFS file open mode
* CIFS file hosting volume
* CIFS share name
* CIFS file path
* Continuously available protection level

 You can specify additional parameters to display only information that
matches those parameters. For example, to display information only about
CIFS files opened on connection ID 2012, run the command with the `-
```

```
connection-id parameter set to ` 2012.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify this parameter, the command only displays the fields that
you specify.

| ``[-instance  ]`` }::
If you specify this parameter, the command displays detailed information
about matching open CIFS files.

``[-node {<nodename>|local}]`` - Node::
If you specify this parameter, the command displays information about the
open CIFS files on the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information about open
CIFS files on the specified CIFS-enabled Vserver.

``[-file-id <integer>]`` - File ID::
If you specify this parameter, the command displays information about the
open CIFS file that match the specified file ID.

``[-connection-id <integer>]`` - Connection ID::
If you specify this parameter, the command displays information about open
CIFS files that are opened on the specified connection ID.

``[-session-id <integer>]`` - Session ID::
If you specify this parameter, the command displays information about the
CIFS file that are opened on the specified session ID.

``[-connection-count <integer>]`` - Connection Count::
If you specify this parameter, the command displays information about CIFS
files opened through a session that have the specified number of CIFS
connections.

``[-file-type <CIFS File Type>]`` - File Type::
If you specify this parameter, the command displays information about
opened CIFS files that are of the specified file type. The file type can
be any of these: Regular, Symlink, Stream, or Directory.

``[-open-mode <CIFS Open Mode>]`` - Open Mode::
If you specify this parameter, the command displays information about CIFS
files that are opened with the specified mode. The open mode can include
one or more of the following:
+
```

```
  * R - This property specifies that the file is opened for read.
  * W - This property specifies that the file is opened for write.
  * D - This property specifies that the file is opened for delete.

+
The open mode can have multiple values specified as a list with no commas.

``[-hosting-aggregate <aggregate name>]`` - Aggregate Hosting File::
If you specify this parameter, the command displays information about open
CIFS files that reside on the specified aggregate.

``[-hosting-volume <volume name>]`` - Volume Hosting File::
If you specify this parameter, the command displays information about open
CIFS files that reside on the specified volume.

``[-share <Share>]`` - CIFS Share::
If you specify this parameter, the command displays information about CIFS
files that are opened over the specified CIFS share.

``[-path <text>]`` - Path from CIFS Share::
If you specify this parameter, the command displays information about open
CIFS files that match the specified CIFS file path.

``[-share-mode <CIFS Open Mode>]`` - Share Mode::
If you specify this parameter, the command displays information about open
CIFS files that are opened with the specified share mode. The share mode
can include one or more of the following:
+
  * R - This property specifies that the file is shared for read.
  * W - This property specifies that the file is shared for write.
  * D - This property specifies that the file is shared for delete.

+
The share mode can have multiple values specified as a list with no
commas.

``[-range-locks <integer>]`` - Range Locks::
If you specify this parameter, the command displays information about open
CIFS files that have the specified number of range locks.

``[-continuously-available <CIFS Open File Protection>]`` - Continuously
Available::
If you specify this parameter, the command displays information about open
CIFS files with or without continuously available protection. The open
files are "continuously available" if they are opened from an SMB 3 client
through a share with the "continuously_available" property set. These open
```

```
files are capable of non-disruptively recovering from takeover and
giveback as well as general aggregate relocation between partners in a
high-availability relationship. Streams opened through a continuously
available share are permitted, but are not currently made continuously
available. Directories may be opened through a continuously available
share, but, by design, will not appear continuously available as clients
do not open them that way. These protection levels are applicable to the
files on read/write volumes residing on storage failover aggregates.
+
The continuously available status can be one of the following:
+

+
  * No - The open file is not continuously available.
  * Yes - The open file is continuously available.



``[-reconnected <text>]`` - Reconnected::
If you specify this parameter, the command displays information about open
CIFS files that have the specified reconnected state. The reconnected
state can be one of the following:
+
  * No - The open file is not reconnected.
  * Yes - The open file is reconnected.



``[-flexgroup-msid <integer>]`` - FlexGroup MSID::
If you specify this parameter, the command displays information about open
CIFS files that reside on the volume within the FlexGroup with the
specified MSID..

== Examples

The following example displays information about all open CIFS files:
```

cluster1::> vserver cifs session file show

Node: node1
Vserver: vs1
Connection: 2192
Session: 1
Connection Count: 4
File File Open Hosting Continuously
ID Type Mode Volume Share Available
------- --------- ---- --------------- --------------------- ------------
7 Regular rw rootvs1 rootca Yes

Path: \win8b8.txt

```
The following example displays information about a CIFS file with file-id
7.
```

cluster1::> vserver cifs session file show -file-id 7 -instance
Node: node1
Vserver: vs1
File ID: 7
Connection ID: 2192
Session ID: 1
Connection count: 4
File Type: Regular
Open Mode: rw
Aggregate Hosting File: aggr1
Volume Hosting File: rootvs1
CIFS Share: rootca
Path from CIFS Share: \win8b8.txt
Share Mode: rd
Range Locks: 0
Continuously Available: Yes
Reconnected: No

```
[[ID9be3430bab76806612772b5b71fe7390]]
= vserver cifs share create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a CIFS share

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share create` command creates a CIFS share.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver on which you want to
create a CIFS share.

``-share-name <Share>`` - Share::
This parameter specifies the name of the CIFS share that you want to
```

```
create. A share name can be up to 80 characters long. If this is a home
directory share (designated as such by specifying the ``_homedirectory_``
on the `-share-properties` parameter), you can include %w (Windows user
name), %u (UNIX user name) and %d (Windows domain name) variables in any
combination with this parameter to generate shares dynamically, with the
resultant share names based on the authenticating user's Windows user
name, UNIX user name, and/or Windows domain name. If the share is used by
administrators to connect to other users' home directory (the option is-
home-dirs-access-for-admin-enabled is set to true) or by a user to connect
to other users' home directory (the option is-home-dirs-access-for-public-
enabled is set to true) , the dynamic share pattern must be preceded by a
tilde (~).

``-path <text>`` - Path::
This parameter specifies the path to the CIFS share. This path must exist
in a volume. A directory path name can be up to 256 characters long. If
there is a space in the path name, you must enclose the entire string in
quotation marks (for example, "/new volume/mount here"). If this is a home
directory share as specified by value of home directory on the `-share-
properties` parameter, you can make the path name dynamic by specifying
the %w (Windows user name), %u (UNIX user name), or %d (domain name)
variables or any of their combination as a part of the value of this
parameter.

``[-share-properties <share properties>,...]`` - Share Properties::
This optional parameter specifies a list of properties for the share. The
list can include one or more of the following:
+
  * homedirectory - This property specifies that the share and path names
are dynamic. Specify this value for a home directory share. In a home
directory share, ONTAP can dynamically generate the share's name and path
by substituting %w, %u, and %d variables with the corresponding Windows
user name, UNIX user name, and domain, respectively, specified as the
value of the `-share-name` and `-path` parameters. For instance, if a
dynamic share is defined with a name of ``_%d_%w_`` , a user logged on as
barbara from a domain named ``_FIN_`` sees the share as ``_FIN_barbara_``
. Using the ``_homedirectory_`` value specifies that the share and path
names are dynamically expanded. This property cannot be added or removed
after share creation.
  * oplocks - This property specifies that the share uses opportunistic
locks, also known as client-side caching. Oplocks are enabled on shares by
default; however, some applications do not work well when oplocks are
enabled. In particular, database applications such as Microsoft Access are
vulnerable to corruption when oplocks are enabled. An advantage of shares
is that a single path can be shared multiple times, with each share having
different properties. For instance, if a path named /dept/finance contains
```

```
both a database and other types of files, you can create two shares to it,
one with oplocks disabled for safe database access and one with oplocks
enabled for client-side caching.
  * browsable - This property allows Windows clients to browse the share.
This is the default initial property for all shares.
  * showsnapshot - This property specifies that snapshots can be viewed
and traversed by clients.
  * changenotify - This property specifies that the share supports
ChangeNotify requests. This is a default initial property for all shares.
  * attributecache - This property enables the file attribute caching on
the CIFS share in order to provide faster access of attributes over SMB
1.0.
+
NOTE: For certain workloads, stale file attribute data could be delivered
to a client.
+

  * continuously-available - This property permits SMB clients that
support it to open files in a persistent manner. Files opened this way are
protected from disruptive events, such as failover and giveback. This
option is not supported for FlexGroup volumes or workgroup CIFS servers.
  * branchcache - This property specifies that the share allows clients to
request BranchCache hashes on the files within this share. This option is
useful only if you specify ``_per-share_`` as the operating mode in the
CIFS BranchCache configuration, and also specify the ``_"oplocks"_`` share
property.
  * access-based-enumeration - This property specifies that Access Based
Enumeration is enabled on this share. ABE-filtered shared folders are
visible to a user based on that individual user's access rights,
preventing the display of folders or other shared resources that the user
does not have rights to access.
  * namespace-caching - This property specifies that the SMB clients
connecting to this share can cache the directory enumeration results
returned by the CIFS servers.
  * encrypt-data - This property specifies that SMB encryption must be
used when accessing this share. Clients that do not support encryption
will not be able to access this share.
  * show-previous-versions - This property specifies that the previous
version can be viewed and restored from the client. This property is
enabled by default.



``[-symlink-properties {enable|hide|read-only|symlinks|symlinks-and-
widelinks|disable|no-strict-security}]`` - Symlink Properties::
This optional parameter specifies how the storage system presents UNIX
```

```
symbolic links (symlinks) to CIFS clients. The default value for this
parameter is "symlinks". The list can include one or more of the
following:
+
  * enable (DEPRECATED*) - This property enables both local symlinks and
wide links for read-write access. DFS advertisements are generated for
both local symlinks and wide links even if the CIFS option `-is-advertise-
dfs-enabled` is set to false.
  * hide (DEPRECATED*) - This property hides symlinks. DFS advertisements
are generated if the CIFS option `-is-advertise-dfs-enabled` is set to
true.
  * read-only (DEPRECATED*) - This property enables symlinks for read-only
access.
  * symlinks - This property enables local symlinks for read-write access.
DFS advertisements are not generated even if the CIFS option `-is-
advertise-dfs-enabled` is set to true.
  * symlinks-and-widelinks – This property enables both local symlinks and
wide links for read-write access. DFS advertisements are generated for
both local symlinks and wide links even if the CIFS option `-is-advertise-
dfs-enabled` is set to false.
  * disable - This property disables symlinks and wide links. DFS
advertisements are not generated even if the CIFS option `-is-advertise-
dfs-enabled` is set to true.
  * no-strict-security - This property enables clients to follow symlinks
outside share boundaries.


+
NOTE: * The ``_enable_`` , ``_hide_`` , and ``_read-only_`` parameters are
deprecated and may be removed in a future release of ONTAP.
+

+
NOTE: The no_strict_security setting does not apply to wide links.
+


``[-file-umask <Octal Integer>]`` - File Mode Creation Mask::
This optional parameter specifies the default UNIX umask for new files
created on the share.

``[-dir-umask <Octal Integer>]`` - Directory Mode Creation Mask::
This optional parameter specifies the default UNIX umask for new
directories created on the share.

``[-comment <text>]`` - Share Comment::
```

```
This optional parameter specifies a text comment for the share that is
made available to Windows clients. The comment can be up to 256 characters
long. If there is a space in the descriptive remark or the path, you must
enclose the entire string in quotation marks (for example, "This is
engineering's share.").

``[-attribute-cache-ttl
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - File Attribute
Cache Lifetime::
This optional parameter specifies the lifetime for the attribute cache
share property, which you specify as the value of the -share-properties
parameter.
+
NOTE: This value is useful only if you specify attributecache as a share
property.
+



``[-offline-files {none|manual|documents|programs}]`` - Offline Files::
This optional parameter allows Windows clients to cache data on this
share. The actual caching behavior depends upon the Windows client. The
value can be one of the following:
+
  * none - Disallows Windows clients from caching any files on this share.
  * manual - Allows users on Windows clients to manually select files to
be cached.
  * documents - Allows Windows clients to cache user documents that are
used by the user for offline access.
  * programs - Allows Windows clients to cache programs that are used by
the user for offline access and may use those files in an offline mode
even if the share is available.



``[-vscan-fileop-profile {no-scan|standard|strict|writes-only}]`` - Vscan
File-Operations Profile::
This optional parameter controls which operations trigger virus scans. The
value can be one of the following:
+
  * no-scan: Virus scans are never triggered for this share.
  * standard: Virus scans can be triggered by open, close, and rename
operations. This is the default profile.
  * strict: Virus scans can be triggered by open, read, close, and rename
operations.
  * writes-only: Virus scans can be triggered only when a file that has
been modified is closed.
```

```
``[-max-connections-per-share <integer>]`` - Maximum Tree Connections on
Share::
This optional parameter specifies the maximum number of simultaneous
connections on the new share. This limit is at the node level, not the
Vserver or cluster level. The default for this parameter is 4294967295.
The value 4294967295 indicates no limit. The allowed range for this
parameter is (1 through 4294967295).

``[-force-group-for-create <text>]`` - UNIX Group for File Create::
This optional parameter specifies that all files that CIFS users create in
a specific share belong to the same group (also called the "force-group").
The "force-group" must be a predefined group in the UNIX group database.
This setting has no effect unless the security style of the volume is UNIX
or mixed security style. If "force-group" has been specified for a share,
the following becomes true for the share:
+
  * Primary GID of the CIFS users who access this share is temporarily
changed to the GID of the "force-group".
  * All files in this share that CIFS users create belong to the same
"force-group", regardless of the primary GID of the file owner.



== Examples

The following example creates a CIFS share named SALES_SHARE on a Vserver
named vs1. The path to the share is /sales.
```

cluster1::> vserver cifs share create -vserver vs1 -share-name SALES_SHARE -path /sales -symlink -properties enable

```
 The following example creates a CIFS share named SALES_SHARE on a Vserver
named vs1. The path to the share is /sales and the share uses
opportunistic locks (client-side caching), the share can be browsed by
Windows clients, and a notification is generated when a change occurs.
```

cluster1::> vserver cifs share create -vserver vs1 -share-name SALE -share-properties browsable,changenotify,oplocks, show-previous-versions

```
 The following example creates a CIFS share named DOCUMENTS on a Vserver
named vs1. The path to the share is /documents and the share uses
opportunistic locks (client-side caching), a notification is generated
when a change occurs, and the share allows clients to ask for BranchCache
hashes for files in the share.
```

cluster1::> vserver cifs share create -vserver vs1 -share-name DOCUMENTS path /documents -share -properties branchcache,changenotify,oplocks

```
 The following example creates a CIFS share named DOCUMENTS on a Vserver
named vs1. The path to the share is /documents and the share uses
opportunistic locks (client-side caching), a notification is generated
when a change occurs, and the share allows clients to cache (client-side
caching) user documents on this share.
```

cluster1::> vserver cifs share create -vserver vs1 -share-name DOCUMENTS -path /documents -share -properties changenotify,oplocks -offline-files documents

```
The following example creates a home directory share on a Vserver named
vs1. The path to the share has a %d and %w combination.
```

cluster1::> vserver cifs share create -share-name %d%w -path %d/%w -share-properties homedirectory -vserver vs1

```
The following example creates a home directory share on a Vserver vs1 to
be used with the home directory option s is-home-dirs-access-for-admin-
enabled and/or is-home-dirs-access-for-public-enabled . The path to the
share has a %d and %w combination.
```

cluster1::> vserver cifs share create -share-name %d%w -path %d/%w -share-properties homedirectory -vserver vs1

```
[[ID393e787828782ee6c5cbeb0eb88bb514]]
= vserver cifs share delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a CIFS share

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share delete` command deletes a CIFS share.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver from which you want to delete a CIFS
share.

``-share-name <Share>`` - Share::
This parameter specifies the name of the CIFS share you want to delete.

== Examples

The following example deletes a CIFS share named share1 from a Vserver
named vs1.
```

cluster1::> vserver cifs share delete -vserver vs1 -share-name share1

```
[[ID093988db4bb367ff8e04d0cf45e16dea]]
= vserver cifs share modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a CIFS share

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver cifs share modify` command modifies a CIFS share.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the CIFS-enabled Vserver containing the CIFS
share you want to modify.

``-share-name <Share>`` - Share::
This parameter specifies the name of the CIFS share that you want to
create. A share name can be up to 80 characters long. If this is a home
directory share (designated as such by specifying the ``_homedirectory_``
on the `-share-properties` parameter), you can include %w (Windows user
name), %u (UNIX user name) and %d (Windows domain name) variables in any
combination with this parameter to generate shares dynamically, with the
resultant share names based on the authenticating user.s Windows user
name, UNIX user name, and/or Windows domain name.

``[-path <text>]`` - Path::
This parameter specifies the path to the CIFS share. This path must exist
in a volume. A directory path name can be up to 256 characters long. If
there is a space in the path name, you must enclose the entire string in
quotation marks (for example, "/new volume/mount here"). If this is a
``_homedirectory_`` share as specified by value of home directory on the
`-share-properties` parameter, a dynamic path name must be specified using
%w (Windows user name), %u (UNIX user name), or %d (domain name) variables
or any of their combination as a part of the value of this parameter. If
this is a ``_continuously-available_`` share as specified by value of
continuously-available on the `-share-properties` parameter, the path must
not be within a FlexGroup because this property is not supported for
FlexGroup volumes.

``[-symlink-properties {enable|hide|read-only|symlinks|symlinks-and-
widelinks|disable|no-strict-security}]`` - Symlink Properties::
This optional parameter specifies how the storage system presents UNIX
symbolic links (symlinks) to CIFS clients. The list can include one or
more of the following:
+
  * enable (DEPRECATED*) - This property enables both local symlinks and
wide links for read-write access. DFS advertisements are generated for
both local symlinks and wide links even if the CIFS option `-is-advertise-
dfs-enabled` is set to false.
  * hide (DEPRECATED*) - This property hides symlinks. DFS advertisements
```

```
are generated if the CIFS option `-is-advertise-dfs-enabled` is set to
true.
  * read-only (DEPRECATED*) - This property enables symlinks for read-only
access.
  * symlinks - This property enables local symlinks for read-write access.
DFS advertisements are not generated even if the CIFS option `-is-
advertise-dfs-enabled` is set to true.
  * symlinks-and-widelinks – This property enables both local symlinks and
wide links for read-write access. DFS advertisements are generated for
both local symlinks and wide links even if the CIFS option `-is-advertise-
dfs-enabled` is set to false.
  * disable - This property disables symlinks and wide links. DFS
advertisements are not generated even if the CIFS option `-is-advertise-
dfs-enabled` is set to true.
  * no-strict-security - This property enables clients to follow symlinks
outside share boundaries.



+
NOTE: The read_only setting does not apply to wide links.
+

+
NOTE: * The ``_enable_`` , ``_hide_`` , and ``_read-only_`` parameters are
deprecated and may be removed in a future release of ONTAP.
+

+
NOTE: The no_strict_security setting does not apply to wide links.
+



``[-file-umask <Octal Integer>]`` - File Mode Creation Mask::
This optional parameter specifies the default UNIX umask for new files
created on the share.

``[-dir-umask <Octal Integer>]`` - Directory Mode Creation Mask::
This optional parameter specifies the default UNIX umask for new
directories created on the share.

``[-comment <text>]`` - Share Comment::
This optional parameter specifies a text comment for the share that is
made available to Windows clients. The comment can be up to 256 characters
long. If there is a space in the descriptive remark or the path, you must
enclose the entire string in quotation marks (for example, "This is
engineering's share.").
```

```
``[-attribute-cache-ttl
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - File Attribute
Cache Lifetime::
This optional parameter specifies the lifetime for the attribute cache
share property, which you specify as the value of the -share-properties
parameter.
+
NOTE:  This value is useful only if you specify attributecache as a share
property.
+


``[-offline-files {none|manual|documents|programs}]`` - Offline Files::
This optional parameter allows Windows clients to cache data on this
share.The actual caching behavior depends upon the Windows client. The
value can be one of the following:
+
  * none - Disallows Windows clients from caching any files on this share.
  * manual - Allows users on Windows clients to manually select files to
be cached.
  * documents - Allows Windows clients to cache user documents that are
used by the user for offline access.
  * programs - Allows Windows clients to cache programs that are used by
the user for offline access and may use those files in an offline mode
even if the share is available.



``[-vscan-fileop-profile {no-scan|standard|strict|writes-only}]`` - Vscan
File-Operations Profile::
This optional parameter controls which operations trigger virus scans. The
value can be one of the following:
+
  * no-scan: Virus scans are never triggered for this share.
  * standard: Virus scans can be triggered by open, close, and rename
operations. This is the default profile.
  * strict: Virus scans can be triggered by open, read, close, and rename
operations.
  * writes-only: Virus scans can be triggered only when a file that has
been modified is closed.



``[-max-connections-per-share <integer>]`` - Maximum Tree Connections on
Share::
This optional parameter specifies a maximum number of simultaneous
```

```
connections to the share. This limit is at the node level, not the Vserver
or cluster level. The default for this parameter is 4294967295. The value
4294967295 indicates no limit. The allowed range for this parameter is (1
through 4294967295).

``[-force-group-for-create <text>]`` - UNIX Group for File Create::
This optional parameter specifies that all files that CIFS users create in
a specific share belong to the same group (also called the "force-group").
The "force-group" must be a predefined group in the UNIX group database.
This setting has no effect unless the security style of the volume is UNIX
or mixed security style. You can disable this option by passing a null
string "".

== Examples

The following example modifies a CIFS share named SALES_SHARE on a Vserver
named vs1. The share uses opportunistic locks. The file mask is set to 644
and the directory mask to 777.
```

cluster1::> vserver cifs share modify -vserver vs1 -share-name SALES_SHARE -symlink-properties hide -file -umask 644 -dir-umask 777

```
The following example modifies a CIFS share named DOCUMENTS on a Vserver
named vs1. This triggers a virus scan on write-only files in the share.
```

cluster1::> vserver cifs share modify -vserver vs1 -share-name DOCUMENTS -vscan-fileop-profile writes-only

```
The following example modifies a CIFS share named DOCUMENTS on a Vserver
vs1. The share allows client to cache (client-side caching) user documents
on this share.
```

cluster1::> vserver cifs share modify -vserver vs1 -share-name DOCUMENTS -offline-files documents

```
The following example modifies a CIFS share named DOCUMENTS on a Vserver
named vs1. The optional parameter "force-group-for-create" can be disabled
by passing the null string as parameter to "force-group-for-create"
option.
```

cluster1::> cifs share modify -vserver vs1 -share-name DOCUMENTS -force-group-for-create ""

```
The following example modifies the symlink property of all the shares on
all the Vserver to "enable".
```

cluster1::> vserver cifs share modify -vserver * -share-name * -symlink-properties enable

```
[[ID4de865841c57aa56cbebf829dafab7af]]
= vserver cifs share show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display CIFS shares

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share show` command displays information about CIFS
shares. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information about all CIFS shares:

* Vserver name
* CIFS share name
* Path
* Share properties
* Comment

 You can specify additional parameters to display only information that
matches those parameters. For example, to display information only about
CIFS shares that use dynamic shares, run the command with the `-share-
properties dynamicshare` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify this parameter, the command only displays the fields that
you specify.

| ``[-shadowcopy  ]``::
If you specify this parameter, the command displays information only about
CIFS shadow copy shares.

| ``[-umask  ]``::
If you specify this parameter, the command displays file and directory
masks for CIFS shares.
```

```
| ``[-instance  ]`` }::
If you specify this parameter, the command displays detailed information
about all CIFS shares.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
CIFS shares on the specified CIFS-enabled Vserver.

``[-share-name <Share>]`` - Share::
If you specify this parameter, the command displays information only about
the CIFS share or shares that match the specified name.

``[-cifs-server <NetBIOS>]`` - CIFS Server NetBIOS Name::
If you specify this parameter, the command displays information only about
the CIFS share or shares that use the CIFS-enabled Vserver with the
specified CIFS server name.

``[-path <text>]`` - Path::
If you specify this parameter, the command displays information only about
the CIFS share or shares that have the specified path.

``[-share-properties <share properties>,...]`` - Share Properties::
If you specify this parameter, the command displays information only about
the CIFS share or shares that have the specified share properties.

``[-symlink-properties {enable|hide|read-only|symlinks|symlinks-and-
widelinks|disable|no-strict-security}]`` - Symlink Properties::
If you specify this parameter, the command displays information only about
the CIFS share or shares that have the specified symbolic link properties.

``[-file-umask <Octal Integer>]`` - File Mode Creation Mask::
If you specify this parameter, the command displays information only about
the CIFS share or shares that use the specified file mask.

``[-dir-umask <Octal Integer>]`` - Directory Mode Creation Mask::
If you specify this parameter, the command displays information only about
the CIFS share or shares that use the specified directory mask.

``[-comment <text>]`` - Share Comment::
If you specify this parameter, the command displays information only about
the CIFS share or shares that have the specified comment.

``[-acl <text>,...]`` - Share ACL::
If you specify this parameter, the command displays information only about
the CIFS share or shares that have the specified ACL.
```

```
``[-attribute-cache-ttl
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - File Attribute
Cache Lifetime::
If you specify this parameter, the command displays information only about
the CIFS share or shares that have the specified attribute-cache-ttl for
attribute cache.

``[-volume <volume name>]`` - Volume Name::
If you specify this parameter, the command displays information only about
the CIFS shares that are present in this volume.

``[-offline-files {none|manual|documents|programs}]`` - Offline Files::
If you specify this parameter, the command displays information only about
the CIFS shares that have the specified Offline Files properties.

``[-vscan-fileop-profile {no-scan|standard|strict|writes-only}]`` - Vscan
File-Operations Profile::
If you specify this parameter, the command displays information only about
the CIFS shares that have the specified Vscan fileop profile.

``[-max-connections-per-share <integer>]`` - Maximum Tree Connections on
Share::
If you specify this parameter, the command displays information only about
the CIFS shares that have the specified maximum connections per share
configured.

``[-force-group-for-create <text>]`` - UNIX Group for File Create::
This optional parameter displays information about the CIFS shares that
have the specified "force-group" parameter configured.

== Examples

The following example displays information about all CIFS shares:
```

cluster1::> vserver cifs share show
Vserver Share Path Properties Comment ACL
--- --- --- --- --- ---
vs1 ROOTSHARE / oplocks Share CNC \
browsable mapped Everyone /
changenoti to top Full
fy of Control
Vserver
global
namespac
e
vs1 admin$ / browsable - -
vs1 c$ / oplocks - BUILTIN\Admi
browsable nistrators /

changenoti Full
fy Control
vs1 ipc$ / browsable - -
4 entries were displayed.

```
The following example displays information about a CIFS share named
SALES_SHARE on a Vserver named vs1.
```

cluster1::> vserver cifs share show -vserver vs1 -share-name SALES_SHARE
Vserver: vs1
Share: SALES_SHARE
CIFS Server NetBIOS Name: WINDATA
Path: /sales
Share Properties: oplocks
browsable
Symlink Properties: enable
File Mode Creation Mask: -
Directory Mode Creation Mask: -
Share Comment: -
Share ACL: Everyone / Full Control
File Attribute Cache Lifetime: -
Offline Files: manual
Vscan File-Operations Profile: standard

```
[[ID3019986f194e500b5046a439c2d666c8]]
= vserver cifs share access-control create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an access control list

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share access-control create` command adds a user or
group to a CIFS share's ACL.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the CIFS
share.
```

```
``-share <Share>`` - Share Name::
This parameter specifies the name of the CIFS share.

``-user-or-group <TextNoCase>`` - User/Group Name::
This parameter specifies the user or group to add to the CIFS share's
access control list. If you specify the user name, you must include the
user's domain using the format "domain\username". The user-or-group
parameter is case-insensitive text.

``[-user-group-type {windows|unix-user|unix-group}]`` - User or Group
Type::
This parameter specifies the type of the user or group to add to the CIFS
share's access control list. The default type is windows. The user-group-
type can be one of the following:
+
  * windows
  * unix-user
  * unix-group



``-permission <access rights>`` - Access Type::
This parameter specifies the permissions for the user or group. The
permissions can be one of the following:
+
  * No_access
  * Read
  * Change
  * Full_Control



== Examples

The following example adds the windows group "Everyone" with
"Full_Control" permission to the access control list of the share "vol3".
```

vs1::> vserver cifs share access-control create -share vol3 -user-or-group Everyone -user-group-type windows -permission Full_Control

```
The following example adds the unix-user "pcuser" and unix-group "daemon"
with "read" permission to the access control list of the share "vol3".
```

vs1::> vserver cifs share access-control create -share vol3 -user-or-group pcuser -user-group-type unix-user -permission read
vs1::> vserver cifs share access-control create -share vol3 -user-or-group daemon -user-group-type unix-group

-permission read

```
[[IDa253f907d64f5f74013b34d2a91b4a95]]
= vserver cifs share access-control delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an access control list

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share access-control delete` command deletes a user or
group from a CIFS share's ACL.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the CIFS
share.

``-share <Share>`` - Share Name::
This parameter specifies the name of the CIFS share.

``-user-or-group <TextNoCase>`` - User/Group Name::
This parameter specifies the user or group to delete from the CIFS share's
access control list. If you specify a user name, you must include the
user's domain using the format "domain\username". The user-or-group
parameter is case-insensitive text.

``[-user-group-type {windows|unix-user|unix-group}]`` - User or Group
Type::
This parameter specifies the type of the user or group to delete from the
CIFS share's access control list. The default type is windows. The user-
group-type can be one of the following:
+
  * windows
  * unix-user
  * unix-group
```

```
== Examples

The following example deletes the group "Everyone" for the access control
list of share "vol3".
```

vs1::> vserver cifs share access-control delete -share vol3 -user-or-group Everyone

```
[[ID585ce326a8fd41317c0c116c6ac7ab46]]
= vserver cifs share access-control modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an access control list

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share access-control modify` command modifies the
permissions of a user or group in a CIFS share's ACL.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the CIFS share
whose ACL you want to modify.

``-share <Share>`` - Share Name::
This parameter specifies the name of the CIFS share whose ACL you want to
modify.

``-user-or-group <TextNoCase>`` - User/Group Name::
This parameter specifies the user or group to modify. If you specify the
user name, you must include the user's domain using the format
"domain\username". The user-or-group parameter is case-insensitive text.

``[-user-group-type {windows|unix-user|unix-group}]`` - User or Group
Type::
This parameter specifies the type of the user or group to modify. The
default type is windows. The user-group-type can be one of the following:
+
```

```
  * windows
  * unix-user
  * unix-group



``[-permission <access rights>]`` - Access Type::
This parameter specifies the permissions for the user or group. The
permissions can be one of the following:
+
  * No_access
  * Read
  * Change
  * Full_Control



== Examples

The following example modifies the access control list for a share named
"vol3". It changes the permission for the windows group "Everyone" to
"Full_Control".
```

vs1::*> vserver cifs share access-control modify -share vol3 -user-or-group Everyone -user-group-type windows -permission Full_Control

```
The following example modifies the access control list for a share named
"vol3". It changes the permission for the unix-user "pcuser" and unix-
group "daemon" to "change".
```

vs1::> vserver cifs share access-control modify -share vol3 -user-or-group pcuser -user-group-type unix-user -permission change
vs1::> vserver cifs share access-control modify -share vol3 -user-or-group daemon -user-group-type unix-group -permission change

```
[[ID01c599501e6f8f6749b5a885f157aa8c]]
= vserver cifs share access-control show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display access control lists on CIFS shares
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share access-control show` command displays the ACLs of
CIFS shares.

== 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.

``[-vserver <vserver name>]`` - Vserver::
This optional parameter specifies the name of the Vserver containing the
share for which you want to display the access control list.

``[-share <Share>]`` - Share Name::
This optional parameter specifies the name of the CIFS share for which you
want to display the access control list.

``[-user-or-group <TextNoCase>]`` - User/Group Name::
If you specify this optional parameter, the command displays only access
control lists for the CIFS shares that have ACLs matching the specified
user or group.

``[-user-group-type {windows|unix-user|unix-group}]`` - User or Group
Type::
If you specify this optional parameter, the command displays only access
control lists for the CIFS shares that have ACLs matching the specified
user-group-type. The user-group-type can be one of the following:
+
  * windows
  * unix-user
  * unix-group



``[-permission <access rights>]`` - Access Type::
 If you specify this optional parameter, the command displays only access
control lists for the CIFS shares that have ACLs matching the specified
```

```
permission. The permissions can be one of the following:
+
  * No_access
  * Read
  * Change
  * Full_Control



``[-winsid <windows sid>]`` - Windows SID::
 If you specify this optional parameter, the command displays only access
control lists for the CIFS shares that have ACLs matching the specified
Windows SID.

``[-unixId <integer>]`` - User/Group ID::
 If you specify this optional parameter, the command displays only access
control lists for the CIFS shares that have ACLs matching the specified
unix user/group Identifier.

``[-access-mask <Hex Integer>]`` - Access mask::
If you specify this optional parameter, the command displays only access
control lists for the CIFS shares that have ACLs matching the specified
access rights.

== Examples

The following example displays all the ACLs for shares in Vserver vs1.
```

vs1::> vserver cifs share access-control show
Share User/Group User/Group Access
Vserver Name Name Type Permission
--- --- --- --- ---
vs1 vol3 CIFSQA\administrator windows Read
vs1 vol3 Everyone windows Full_Control
vs1 vol3 pcuser unix-user Read
vs1 vol3 daemon unix-group Read

```
[[IDe4fc7fa28b5cbad37cf6875208cb7421]]
= vserver cifs share properties add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to the list of share properties
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share properties add` command adds share properties to
the list of share properties of an existing CIFS share. You can add one or
more share properties. You can add additional share properties at any time
by rerunning this command. Any share properties that you have previously
specified will remain in effect and newly added properties are appended to
the existing list of share properties.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the CIFS share
whose share properties you want to add.

``-share-name <Share>`` - Share::
This parameter specifies the name of the CIFS share.

``-share-properties <share properties>,...`` - Share Properties::
This parameter specifies the list of share properties you want to add to
the CIFS share. The share properties can be one or more of the following:
+
  * oplocks - This property specifies that the share uses opportunistic
locks, also known as client-side caching. This is a default initial
property for all shares; however, some applications do not work well when
oplocks are enabled. In particular, database applications such as
Microsoft Access are vulnerable to corruption when oplocks are enabled. An
advantage of shares is that a single path can be shared multiple times,
with each share having different properties. For instance, if a path named
``_/dept/finance_`` contains both a database and other types of files, you
can create two shares to it, one with oplocks disabled for safe database
access and one with oplocks enabled for client-side caching.
  * browsable - This property allows Windows clients to browse the share.
This is a default initial property for all shares.
  * showsnapshot - This property specifies that snapshots can be viewed
and traversed by clients.
  * changenotify - This property specifies that the share supports
ChangeNotify requests. This is a default initial property for all shares.
  * attributecache - This property enables the file attribute caching on
the CIFS share in order to provide faster access of attributes over SMB
1.0.
+
NOTE:  For certain workloads, stale file attribute data could be delivered
```

```
to a client.
+

  * continuously-available - This property permits SMB clients that
support it to open files in a persistent manner. Files opened this way are
protected from disruptive events, such as failover and giveback. This
option is not supported for FlexGroup volumes or workgroup CIFS servers.
  * branchcache - This property specifies that the share allows clients to
request BranchCache hashes on the files within this share. This option is
useful only if you specify "per-share" as the operating mode in the CIFS
BranchCache configuration, and also specify the ``_"oplocks"_`` share
property.
  * access-based-enumeration - This property specifies that Access Based
Enumeration(ABE) is enabled on this share. ABE-filtered shared folders are
visible to a user based on that individual user's access rights,
preventing the display of folders or other shared resources that the user
does not have rights to access.
  * namespace-caching - This property specifies that the SMB clients
connecting to this share can cache the directory enumeration results
returned by the CIFS servers.
  * encrypt-data - This property specifies that SMB encryption must be
used when accessing this share. Clients that do not support encryption
will not be able to access this share.
  * show-previous-versions - This property specifies that the previous
version can be viewed and restored from the client. This property is
enabled by default.


+
NOTE: The oplock, browsable, changenotify and show-previous-versions share
properties are assigned to a share by default. If you have removed them
from a share, you can use the vserver cifs share properties add command to
add these properties to the share.
+


== Examples

The following example adds the "showsnapshot" and "changenotify"
properties to a share named "sh1".
```

cluster1::> vserver cifs share properties add -vserver vs1 -share-name sh1 -share-properties showsnapshot,changenotify

```
[[ID4b1f445f05c1b73daab9e485fce002b3]]
= vserver cifs share properties remove
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from the list of share properties

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share properties remove` command removes share
properties from the list of share properties of an existing CIFS share.
You can remove one or more share properties. You can remove additional
share properties at any time by rerunning this command. Any existing share
properties that you do not remove remain in effect.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the CIFS share
whose share properties you want to remove.

``-share-name <Share>`` - Share::
This parameter specifies the name of the CIFS share.

``-share-properties <share properties>,...`` - Share Properties::
This parameter specifies the list of share properties you want to remove
from the CIFS share. The share properties can be one or more of the
following:
+
  * oplocks - This property specifies that the share uses opportunistic
locks, also known as client-side caching. Oplocks are enabled on shares by
default; however, some applications do not work well when oplocks are
enabled. In particular, database applications such as Microsoft Access are
vulnerable to corruption when oplocks are enabled. An advantage of shares
is that a single path can be shared multiple times, with each share having
different properties. For instance, if a path named ``_/dept/finance_``
contains both a database and other types of files, you can create two
shares to it, one with oplocks disabled for safe database access and one
with oplocks enabled for client-side caching.
  * browsable - This property allows Windows clients to browse the share.
  * showsnapshot - This property specifies that snapshots can be viewed
and traversed by clients.
```

```
  * changenotify - This property specifies that the share supports
ChangeNotify requests. This is a default initial property for all shares.
  * attributecache - This property enables the file attribute caching on
the CIFS share in order to provide faster access of attributes over SMB
1.0.
+
NOTE:  For certain workloads, stale file attribute data could be delivered
to a client.
+

  * continuously-available - This property permits SMB clients that
support it to open files in a persistent manner. Files opened this way are
protected from disruptive events, such as failover and giveback. This
option is not supported for FlexGroup volumes or workgroup CIFS servers.
  * branchcache - This property specifies that the share allows clients to
request BranchCache hashes on the files within this share. This option is
useful only if you specify "per-share" as the operating mode in the CIFS
BranchCache configuration, and also specify the ``_"oplocks"_`` share
property.
  * access-based-enumeration - This property specifies that Access Based
Enumeration(ABE) is enabled on this share. ABE-filtered shared folders are
visible to a user based on that individual user's access rights,
preventing the display of folders or other shared resources that the user
does not have rights to access.
  * namespace-caching - This property specifies that the SMB clients
connecting to this share can cache the directory enumeration results
returned by the CIFS servers.
  * encrypt-data - This property specifies that SMB encryption must be
used when accessing this share. Clients that do not support encryption
will not be able to access this share.
  * show-previous-versions - This property specifies that the previous
version can be viewed and restored from the client. This property is
enabled by default.



== Examples

The following example removes "showsnapshot" and "changenotify" properties
to a share named "sh1".
```

cluster1::> vserver cifs share properties remove -vserver vs1 -share-name sh1 -share-properties showsnapshot,changenotify

```
[[ID9a03ad8cdb1de8b0938100120767ce9b]]
= vserver cifs share properties show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display share properties

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs share properties show` command displays the CIFS share
properties.

== 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.

``[-vserver <vserver name>]`` - Vserver::
This optional parameter specifies the name of the Vserver containing the
CIFS share for which you want to display share properties.

``[-share-name <Share>]`` - Share::
If you specify this parameter, the command displays share properties only
for the CIFS share that you specify.

``[-share-properties <share properties>,...]`` - Share Properties::
If you specify this parameter, the command displays share properties only
for CIFS shares using the properties you specify. The share properties can
be one or more of the following:
+
  * homedirectory - This property specifies that the share and path names
are dynamic. Specify this value for a home directory share. In a home
directory share, the share's name and path can be generated by
substituting %w and %d variables with the corresponding user's name and
domain, respectively, specified as the value of the `-share-name` and `-
path` parameters. For instance, if a dynamic share is defined with a name
of ``_%d_%w_`` , a user logged on as ``_barbara_`` from a domain named
```

```
``_FIN_`` sees the share as ``_FIN_barbara_`` . Using the homedirectory
value specifies that the share and path names are dynamically expanded.
  * oplocks - This property specifies that the share uses opportunistic
locks, also known as client-side caching.
  * browsable - This property allows Windows clients to browse the share.
  * showsnapshot - This property specifies that snapshots can be viewed
and traversed by clients.
  * changenotify - This property specifies that the share supports Change
Notify requests.
  * attributecache - This property enables the file attribute caching on
the CIFS share in order to provide faster access of attributes over SMB
1.0.
+
NOTE:  For certain workloads, stale file attribute data could be delivered
to a client.
+

  * continuously-available - This property permits SMB clients that
support it to open files in a persistent manner. Files opened this way are
protected from disruptive events, such as failover and giveback. This
attribute is not supported for FlexGroup volumes and workgroup CIFS
servers.
  * branchcache - This property specifies that the share allows clients to
request BranchCache hashes on the files within this share. This option is
useful only if you specify "per-share" as the operating mode in the CIFS
BranchCache configuration, and also specify the ``_"oplocks"_`` share
property.
  * shadowcopy - This property specifies that the share is pointing to a
shadow copy. This attribute cannot be added nor removed from a share.
  * access-based-enumeration - This property specifies that Access Based
Enumeration is enabled on this share. ABE-filtered shared folders are
visible to a user based on that individual user's access rights,
preventing the display of folders or other shared resources that the user
does not have rights to access.
  * namespace-caching - This property specifies that the SMB clients
connecting to this share can cache the directory enumeration results
returned by the CIFS servers.
  * encrypt-data - This property specifies that SMB encryption must be
used when accessing this share. Clients that do not support encryption
will not be able to access this share.
  * show-previous-versions - This property specifies that the previous
version can be viewed and restored from the client. This property is
enabled by default.
```

```
== Examples

The following example displays share properties for shares in Vserver vs1.
```

cluster1::> vserver cifs share properties show
Vserver Share Properties
--- --- ---
vs1 abc oplocks
browsable
changenotify
show-previous-versions
vs1 admin$ browsable
vs1 ipc$ browsable
vs1 sh1 oplocks
browsable
changenotify
show-previous-versions
4 entries were displayed.

```
[[ID88fe8e4391b690ac5ff176682a1b8af9]]
= vserver cifs superuser create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Adds superuser permissions to a CIFS account

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs superuser create` command elevates the privileges of the
specified domain account in this Vserver to superuser. With superuser
privileges, ONTAP bypasses some of the security checks. This command is
not supported for workgroup CIFS servers.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Vserver name.

``-domain <CIFS domain>`` - Domain [.advaced-priv]#(privilege:
advanced)#::
The domain name of accountname.

``-accountname <CIFS account>`` - User [.advaced-priv]#(privilege:
advanced)#::
The domain account to which you want to give superuser privileges.

== Examples

The following example shows how to elevate ExampleUser in EXAMPLE domain
to superuser for a Vserver vs1.
```

vs1::> vserver cifs superuser create -domain EXAMPLE -accountname ExampleUser -vserver vs1

```
[[ID50c263b3bf1a53d4f70e4f6ada303a6f]]
= vserver cifs superuser delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Deletes superuser permissions from a CIFS account

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs superuser delete` command removes the superuser
privileges for the specified domain account in this Vserver. With
superuser privileges, ONTAP bypasses some of the security checks.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Vserver name.

``-domain <CIFS domain>`` - Domain [.advaced-priv]#(privilege:
advanced)#::
The domain name of accountname.

``-accountname <CIFS account>`` - User [.advaced-priv]#(privilege:
advanced)#::
The domain account name you want to remove superuser privileges.

== Examples

The following example shows how to remove superuser privileges for
ExampleUser in EXAMPLE domain for a Vserver vs1.
```

vs1::> vserver cifs superuser delete -domain EXAMPLE -accountname ExampleUser -vserver vs1

```
[[ID307cba5f3d242e7ce5fd7707e17e52b5]]
= vserver cifs superuser show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display superuser permissions for CIFS accounts

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs superuser show` command displays all account names with
superuser privileges. The command output depends on the parameter or
parameters specified with the command. If you do not specify any
parameters, the command displays the following superuser information for
all CIFS servers:

* Vserver name
* CIFS server NetBIOS name
* Domain
* Account Name



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays superuser information
of only the specified Vservers.

``[-domain <CIFS domain>]`` - Domain [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays superuser information
of only for accounts that are in the specified domain.

``[-accountname <CIFS account>]`` - User [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays superuser information
```

```
of only the CIFS servers with the specified superuser account.

``[-cifs-server <NetBIOS>]`` - CIFS Server NetBIOS Name [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays superuser information
of only the Vservers with specified CIFS server name.

== Examples

The following example displays superuser information of all Vservers.
```

vs1::> vserver cifs superuser show

Vserver CIFS Server Domain Account Name
--- --- --- ---
vs1 SMB_SERVER1 CIFSDOMAIN ADMINISTRATOR
vs2 SMB_SERVER2 CIFSDOMAIN ADMINISTRATOR

```
[[IDa3fc7dae1bc0c74a75f45f2b1eb198e2]]
= vserver cifs symlink create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a symlink path mapping

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs symlink create` command creates a symbolic link mapping
for CIFS. A mapping consists of a Vserver name, a UNIX (NFS) path, a CIFS
share name, and a CIFS path. You can also specify a CIFS server name and
whether the CIFS symbolic link is a local link, a free link (obsolete), or
wide link. A local symbolic link maps to the local CIFS share. A free
symbolic link can map anywhere on the local server. A wide symbolic link
maps to any CIFS share on the network. If the target share is a Home
Directory, then the `-home-directory` field must be set to true for
correct processing.

== Parameters

``-vserver <vserver name>`` - Vserver::
```

```
This parameter specifies the Vserver on which you want to create the
mapping.

``-unix-path <text>`` - UNIX Path::
This parameter specifies the UNIX (NFS) path for the mapping.
+
NOTE: It must begin and end with a forward slash (/).
+


``[-share-name <Share>]`` - CIFS Share::
This parameter specifies the CIFS share for the mapping.

``-cifs-path <TextNoCase>`` - CIFS Path::
This parameter specifies the CIFS path for the mapping. Note that this
value is specified by using a UNIX-style path.
+
NOTE: It must begin and end with a forward slash (/).
+


``[-cifs-server <TextNoCase>]`` - Remote NetBIOS Server Name::
This parameter specifies a new CIFS server DNS name, IP address, or
NetBIOS name for the mapping.

``[-locality {local|widelink}]`` - Local or Wide Symlink::
This parameter specifies whether the CIFS symbolic link is a local link, a
free link (obsolete), or wide link. A local symbolic link maps to the
local CIFS share. A free symbolic link can map anywhere on the local
server. A wide symbolic link maps to any CIFS share on the network. The
default setting is `local` . The free link option is obsolete.

``[-home-directory {true|false}]`` - Home Directory::
This parameter specifies whether the target share is a home directory. The
default value is false.
+
NOTE: This field must be set to true when the target share is a Home
Directory for correct processing.
+


== Examples

The following example creates a symbolic link mapping on a Vserver named
vs1. It has the UNIX path /sales/, the CIFS share name SALES_SHARE, and
the CIFS path /mycompany/sales/.
```

cluster1::> vserver cifs symlink create -vserver vs1
-unix-path /sales/ -share-name SALES_SHARE -cifs-path "/mycompany/sales/"

```
The following example creates a symbolic link mapping on a Vserver named
vs1. It has the UNIX path /example/, the CIFS share name EXAMPLE_SHARE,
the CIFS path /mycompany/example/, the CIFS server IP address, and is a
wide link.
```

cluster1::> vserver cifs symlink create -vserver vs1 -unix-path /example/ -share-name EXAMPLE_SHARE
-cifs-path "/mycompany/example/" -cifs-server CIFSSERVER1 -locality widelink

```
[[ID01a746ea778bdd3e196997e9097c1f3d]]
= vserver cifs symlink delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a symlink path mapping

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs symlink delete` command deletes a symbolic link mapping
for CIFS.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver on which the symbolic link mapping is located.

``-unix-path <text>`` - UNIX Path::
This specifies the UNIX (NFS) path of the mapping that you want to delete.

== Examples

The following example deletes a symbolic link mapping to a UNIX path
/example/ from a Vserver named vs1:
```

cluster1::> vserver cifs symlink delete -vserver vs1 -unix-path /example/

```
[[ID27e1aeec052c41cbd3e2727e0557652d]]
= vserver cifs symlink modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a symlink path mapping

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs symlink modify` command modifies the CIFS share name,
CIFS path, CIFS server name, or locality of a symbolic link mapping. It
can also be used to modify the value of `-home-directory` field.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the symbolic link mapping is
located.

``-unix-path <text>`` - UNIX Path::
This parameter specifies the UNIX (NFS) path of the mapping that you want
to modify.
+
NOTE: It must begin and end with a forward slash (/).
+


``[-share-name <Share>]`` - CIFS Share::
This parameter specifies a new CIFS share name for the mapping.

``[-cifs-path <TextNoCase>]`` - CIFS Path::
This parameter specifies a new CIFS path for the mapping. Note that this
value is specified by using a UNIX-style path.
+
NOTE: It must begin and end with a forward slash (/).
+


``[-cifs-server <TextNoCase>]`` - Remote NetBIOS Server Name::
This parameter specifies a new CIFS server DNS name, IP address, or
NetBIOS name for the mapping.
```

```
``[-locality {local|widelink}]`` - Local or Wide Symlink::
This parameter specifies a new locality for the mapping. A local symbolic
link maps to the local CIFS share. A free symbolic link can map anywhere
on the local server. A wide symbolic link maps to any CIFS share on the
network. The default setting is `local` . The free link option is
obsolete.

``[-home-directory {true|false}]`` - Home Directory::
This parameter specifies whether the new target share is a home directory.
+
NOTE: This field must be set to true when the target share is a Home
Directory for correct processing.
+


== Examples

The following example modifies the symbolic link mapping to a UNIX path
/example/ on a Vserver named vs1. The mapping is modified to use the CIFS
path /mycompany/example/.
```

cluster1::> vserver cifs symlink modify -vserver vs1 -unix-path /example/ -cifs-path "/mycompany/example/"

```
The following example modifies the symbolic link mapping to a UNIX path
/example/ on a Vserver named vs1. The mapping is modified to use the CIFS
share name EXAMPLE_SHARE, the CIFS path /mycompany/example/, on the CIFS
server cifs.example.com, and to be a wide link.
```

cluster1::> vserver cifs symlink modify -vserver vs1 -unix-path /example/ -share-name EXAMPLE_SHARE -cifs
-path "/mycompany/example/" -cifs-server cifs.example.com
-locality widelink

```
[[ID16e8fed6e405d8d2dbbd46dbc8142d59]]
= vserver cifs symlink show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show symlink path mappings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver cifs symlink show` command displays the following information
about symbolic link mappings for CIFS:

* Vserver
* UNIX (NFS) path
* The DNS name, IP address, or NetBIOS name of the CIFS server
* CIFS share name
* CIFS path
* Whether the locality of the CIFS server is a local, free, or wide link.
(A local symbolic link maps to the local CIFS share. A free symbolic link
can map anywhere on the local server. A wide symbolic link maps to any
CIFS share on the network. The free link option is deprecated and may be
removed in a future release of ONTAP.)



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information about
symbolic link mappings on the specified Vserver.

``[-unix-path <text>]`` - UNIX Path::
If you specify this parameter, the command displays information only about
the symbolic link mapping that uses the specified UNIX (NFS) path.

``[-share-name <Share>]`` - CIFS Share::
If you specify this parameter, the command displays information only about
the symbolic link mapping or mappings that use the specified CIFS share.

``[-cifs-path <TextNoCase>]`` - CIFS Path::
If you specify this parameter, the command displays information only about
the symbolic link mapping that uses the specified CIFS path.

``[-cifs-server <TextNoCase>]`` - Remote NetBIOS Server Name::
If you specify this parameter, the command displays information only about
```

```
the symbolic link mapping that uses the specified CIFS server.

``[-locality {local|widelink}]`` - Local or Wide Symlink::
If you specify this parameter, the command displays information only about
the symbolic link mappings that have the specified locality.

``[-home-directory {true|false}]`` - Home Directory::
If you specify this parameter, the command displays information only about
the symbolic link mappings that have the target share as a home directory
(if true) or as a static CIFS share (if false).

== Examples

The following example displays information about all symbolic link
mappings for CIFS:
```

cluster1::> vserver cifs symlink show
Vserver Unix Path CIFS Server CIFS Share CIFS Path Locality
---------- ---------- ------------------ ----------- ------------------ --------
vs1 /hr/ 192.0.2.160 HR_SHARE /mycompany/hr/ widelink
vs1 /sales/ WINDATA SALES_SHARE /mycompany/sales/ local
vs1 /web/ cifs.example.com WEB_SHARE /mycompany/web/ widelink
3 entries were displayed.

```
The following example displays information about all symbolic link
mappings that are wide links:
```

cluster1::> vserver cifs symlink show -locality widelink
Vserver Unix Path CIFS Server CIFS Share CIFS Path Locality
---------- ---------- ------------------ ----------- ---------------- --------
vs1 /hr/ 192.0.2.160 HR_SHARE /mycompany/hr/ widelink
vs1 /web/ cifs.example.com WEB_SHARE /mycompany/web/ widelink
2 entries were displayed.

```
[[ID984aaf66efe9141a05fe6944c8021c6d]]
= vserver cifs users-and-groups remove-stale-records
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete the Stale CIFS local users-and-groups records for the specified
vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver cifs users-and-groups remove-stale-records` command removes
Stale local users and groups entries associated with old CIFS server.

== Parameters

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The command deletes Stale local users and groups entries associated with
the specified Vserver.

== Examples

The following example displays the syntax of the command.
```

cluster1::*> vserver cifs users-and-groups remove-stale-records -vserver vs1

```
[[ID314af09ce34254d9bb87c4371b25ca56]]
= vserver cifs users-and-groups update-names
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Update the names of Active Directory users and groups

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description
```

```
The `vserver cifs users-and-groups update-names` command updates the names
of Active Directory users and groups that are registered in local
databases on the cluster and reports the status of the update operations.
This is done so that objects that were renamed in the Active Directory can
be properly displayed and configured in the local databases.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command will only be performed within
the scope of the Vserver that matches the specified Vserver name.

{ ``[-display-failed-only {true|false}]`` - Display Only Failures
[.advaced-priv]#(privilege: advanced)#::
If you set this parameter to true, the command displays only the Active
Directory users and groups that failed to update. If set to false, the
command displays only the Active Directory users and groups that
successfully updated.

| ``[-suppress-all-output {true|false}]`` - Suppress All Output [.advaced-
priv]#(privilege: advanced)# }::
If you set this parameter to true, the command does not display
information about the status of the updates of Active Directory users and
groups. To display information about the status of the updates, set this
parameter to false or do not specify this parameter in the command.

== Examples

The following example updates the names of Active Directory users and
groups associated with Vserver "vs1". In the last case, there is a
dependent chain of names that needs to be updated.
```

cluster1::*> vserver cifs users-and-groups update-names -vserver vs1
Vserver: vs1
SID: S-1-5-21-123456789-234565432-987654321-12345
Domain: EXAMPLE1
Out-of-date Name: dom_user1
Updated Name: dom_user2
Status: Successfully updated
Vserver: vs1
SID: S-1-5-21-123456789-234565432-987654322-23456
Domain: EXAMPLE2
Out-of-date Name: dom_user1
Updated Name: dom_user2
Status: Successfully updated
Vserver: vs1

SID: S-1-5-21-123456789-234565432-987654321-123456
Domain: EXAMPLE1
Out-of-date Name: dom_user3
Updated Name: dom_user4
Status: Successfully updated; also updated SID "S-1-5-21-123456789-234565432-987654321-123457"
to name "dom_user5"; also updated SID "S-1-5-21-123456789-234565432-987654321-123458"
to name "dom_user6"; also updated SID "S-1-5-21-123456789-234565432-987654321-123459"
to name "dom_user7"; also updated SID "S-1-5-21-123456789-234565432-987654321-123460"
to name "dom_user8"

The command completed successfully. 7 Active Directory objects have been updated.

```
[[ID1106380ddaa6ae935eff1155ef5f3039]]
= vserver cifs users-and-groups local-group add-members
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add members to a local group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group add-members` command adds
members to a local group.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-group-name <CIFS name>`` - Group Name::
This specifies the name of the local group.

``-member-names <CIFS name>,...`` - Names of Users or Active Directory
Groups to be Added::
This specifies the list of local users, Active Directory users, or Active
Directory groups to be added to a particular local group.

== Examples

The following example adds a local user "CIFS_SERVER\loc_usr1" and an
Active Directory group "CIFS_SERVER\dom_grp2" to the local group
"CIFS_SERVER\g1".
```

cluster1::> vserver cifs users-and-groups local-group add-members -vserver vs1 -group-name CIFS_SERVER\g1 -member-names CIFS_SERVER\loc_usr1,AD_DOMAIN\dom_grp2

```
[[ID1df1c7a690635cb46f32ac0026d13b5a]]
= vserver cifs users-and-groups local-group create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a local group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group create` command creates a
local group and optionally sets the description of that local group. The
group name must meet the following criteria:

* The group name length must not exceed 256 characters.
* The group name cannot be terminated by a period.
* The group name cannot include commas.
* The group name cannot include any of the following printable characters:
", /, \, [, ], :, |, <, >, +, =, ;, ?, *, @
* The group name cannot include characters in the ASCII range 1-31, which
are non-printable.



== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-group-name <CIFS name>`` - Group Name::
This specifies the name of the local group.

``[-description <TextNoCase>]`` - Description::
This specifies a description for this local group. If the description
contains a space, enclose the parameter in quotation marks.

== Examples

The following example creates a local group "CIFS_SERVER\g1" associated
with Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-group create -vserver vs1 -group-name CIFS_SERVER\g1

```
[[IDcd5c10c2db46d2ff2e9899ca473010aa]]
= vserver cifs users-and-groups local-group delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a local group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group delete` command deletes a
local group. Removing a local group removes its membership records.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-group-name <CIFS name>`` - Group Name::
This specifies the name of the local group to delete.

== Examples

The following example deletes the local group "CIFS_SERVER\g1" associated
with Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-group delete -vserver vs1 -group-name CIFS_SERVER\g1

```
[[IDa400a0c2e968e71f562f328611546c4b]]
= vserver cifs users-and-groups local-group modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a local group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group modify` command modifies
the description of a local group.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-group-name <CIFS name>`` - Group Name::
This specifies the name of the local group.

``[-description <TextNoCase>]`` - Description::
This specifies a description for this local group. If the description
contains a space, enclose the parameter in quotation marks.

== Examples

The following example modifies the description of the local group
"CIFS_SERVER\g1" associated with Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-group modify -vserver vs1 -group-name CIFS_SERVER\g1 -description "Example Description"

```
[[IDb6b9ccd53cbc0dd94bd927bbc9fbfa59]]
= vserver cifs users-and-groups local-group remove-members
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove members from a local group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group remove-members` command
removes members from a local group.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-group-name <CIFS name>`` - Group Name::
This specifies the name of the local group.

``-member-names <CIFS name>,...`` - Names of Users or Active Directory
Groups to be Removed::
This specifies the list of local users, Active Directory users, or Active
Directory groups to be removed from a particular local group.

== Examples

The following example removes the local users "CIFS_SERVER\u1" and
"CIFS_SERVER\u2" from the local group "CIFS_SERVER\g1".
```

cluster1::> vserver cifs users-and-groups local-group remove-members -vserver vs1 -group-name CIFS_SERVER\g1 -member-names CIFS_SERVER\u1,CIFS_SERVER\u2

```
[[ID50956ffeb27173f12b22821ac0745436]]
= vserver cifs users-and-groups local-group rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a local group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group rename` command renames a
local group. The new group name must remain in the same domain as the old
group name. The new group name must meet the following criteria:

* The group name length must not exceed 256 characters.
* The group name cannot be terminated by a period.
* The group name cannot include commas.
* The group name cannot include any of the following printable characters:
", /, \, [, ], :, |, <, >, +, =, ;, ?, *, @
* The group name cannot include characters in the ASCII range 1-31, which
are non-printable.



== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-group-name <CIFS name>`` - Group Name::
This specifies the local group's name.

``-new-group-name <CIFS name>`` - New Group Name::
This specifies the local group's new name.

== Examples

The following example renames the local group "CIFS_SERVER\g_old" to
"CIFS_SERVER\g_new" on Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-group rename -group-name CIFS_SERVER\g_old -new-group -name CIFS_SERVER\g_new -vserver vs1

```
[[IDa5a884dbc602fb38383acd88b7692e42]]
= vserver cifs users-and-groups local-group show-members
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local groups' members

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group show-members` command
displays members of a local group. The members can be local or Active
Directory users or groups.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays group members of
local groups that match the specified Vserver name.

``[-group-name <CIFS name>]`` - Group Name::
If this parameter is specified, the command displays group members of
local groups that match the specified group name.

``[-member <CIFS name>,...]`` - Member Name::
If this parameter is specified, the command displays group members that
match the specified member name. The name can be that of a local user,
Active Directory user, or Active Directory group.
```

```
== Examples

The following example displays members of local groups associated with
Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-group show-members -vserver vs1
Vserver Group Name Members
--- --- ---
vs1 BUILTIN\Administrators CIFS_SERVER\Administrator
AD_DOMAIN\Domain Admins
AD_DOMAIN\dom_grp1
BUILTIN\Users AD_DOMAIN\Domain Users
AD_DOMAIN\dom_usr1
CIFS_SERVER\g1 CIFS_SERVER\u1
6 entries were displayed.

```
[[ID52725eaa7ce2f2ed09b063afe934bb6]]
= vserver cifs users-and-groups local-group show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local groups

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-group show` command displays
local groups.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays information only
about local groups that match the specified Vserver name.

``[-group-name <CIFS name>]`` - Group Name::
If this parameter is specified, the command displays information only
about local groups that match the specified group name.

``[-description <TextNoCase>]`` - Description::
If this parameter is specified, the command displays information only
about local groups that match the specified description.

== Examples

The following example displays all local groups associated with Vserver
"vs1".
```

```
cluster1::> vserver cifs users-and-groups local-group show -vserver vs1
Vserver Group Name Description
------------- ------------------------------ ----------------------------
vs1 BUILTIN\Administrators Built-in Administrators group
vs1 BUILTIN\Backup Operators Backup Operators group
vs1 BUILTIN\Power Users Restricted administrative privileges
vs1 BUILTIN\Users All users
vs1 CIFS_SERVER\g1
vs1 CIFS_SERVER\g2
6 entries were displayed.
```

```
[[ID6d2d2e69aad5ed4bcdf52923a72048b0]]
= vserver cifs users-and-groups local-user create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a local user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-user create` command creates a
local user and optionally sets the attributes for that local user. The
command prompts for the local user's password. + + + The user name must
meet the following criteria: +

* The user name length must not exceed 20 characters.
* The user name cannot be terminated by a period.
* The user name cannot include commas.
* The user name cannot include any of the following printable characters:
", /, \, [, ], :, |, <, >, +, =, ;, ?, *, @
* The user name cannot include characters in the ASCII range 1-31, which
are non-printable.

The password must meet the following criteria:

* The password must be at least six characters in length.
* The password must not contain user account name.
* The password must contain characters from three of the following four
categories:

  * English uppercase characters (A through Z)
```

```
  * English lowercase characters (a through z)
  * Base 10 digits (0 through 9)
  * Special characters: ~, !, @, #, 0, ^, , *, _, -, +, =, `, \, |, (, ),
[, ], :, ;, ", ', <, >, ,, ., ?, /





== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-name <CIFS name>`` - User Name::
This specifies the user name.

``[-full-name <TextNoCase>]`` - Full Name::
This specifies the user's full name. If the full name contains a space,
enclose the full name within double quotation marks.

``[-description <TextNoCase>]`` - Description::
This specifies a description for this local user. If the description
contains a space, enclose the parameter in quotation marks.

``[-is-account-disabled {true|false}]`` - Is Account Disabled::
This specifies whether the user account is enabled or disabled. Set this
parameter to true to disable the account. Set this parameter to false to
enable the account. If this parameter is not specified, the default is to
enable the user account.

== Examples

The following example creates a local user "CIFS_SERVER\u1" associated
with Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-user create -vserver vs1 -user-name CIFS_SERVER\u1

Enter the password:
Confirm the password:

```
[[ID5739619155c0eb50e1795f74f3cac53e]]
= vserver cifs users-and-groups local-user delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a local user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-user delete` command deletes a
local user. Upon deletion, all membership entries for this local user are
deleted.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-name <CIFS name>`` - User Name::
This specifies the user name.

== Examples

The following example deletes the local user "CIFS_SERVER\u1" associated
with Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-user show-membership
(vserver cifs users-and-groups local-user show-membership)
Vserver User Name Membership
--- --- ---
vs1 CIFS_SERVER\Administrator BUILTIN\Administrators
CIFS_SERVER\u1 CIFS_SERVER\g1
2 entries were displayed.

cluster1::> vserver cifs users-and-groups local-user delete -vserver vs1 -user-name CIFS_SERVER\u1

cluster1::> vserver cifs users-and-groups local-user show-membership
Vserver User Name Membership
--- --- ---
vs1 CIFS_SERVER\Administrator BUILTIN\Administrators

```
[[IDb6437cb06907ea12007f73543bd9725e]]
= vserver cifs users-and-groups local-user modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a local user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-user modify` command modifies
attributes of a local user.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-name <CIFS name>`` - User Name::
This specifies the user name.

``[-full-name <TextNoCase>]`` - Full Name::
This specifies the user's full name. If the full name contains a space in
the name, enclose it within double quotation marks

``[-description <TextNoCase>]`` - Description::
This specifies a description for this local user. If the description
contains a space, enclose the parameter in quotation marks.

``[-is-account-disabled {true|false}]`` - Is Account Disabled::
This specifies if the user account is enabled or disabled. Set this
parameter to true to disable the account. Set this parameter to false to
enable the account.

== Examples

The following example modifies the full name of the local user
"CIFS_SERVER\u1".
```

cluster1::> vserver cifs users-and-groups local-user modify -user-name CIFS_SERVER\u1 -full-name "John Smith" -vserver vs1

```
[[IDed49f363ed25b46ff3baefccbf001539]]
= vserver cifs users-and-groups local-user rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename a local user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-user rename` command renames a
local user. The new user name must remain in the same domain as the old
user name. + The new user name must meet the following criteria:

* The user name length must not exceed 20 characters.
* The user name cannot be terminated by a period.
* The user name cannot include commas.
* The user name cannot include any of the following printable characters:
", /, \, [, ], :, |, <, >, +, =, ;, ?, *, @
* The user name cannot include characters in the ASCII range 1-31, which
are non-printable.



== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-name <CIFS name>`` - User Name::
This specifies the user name.

``-new-user-name <CIFS name>`` - New User Name::
This specifies the new user name.

== Examples

The following example renames the local user "CIFS_SERVER\u_old" to
"CIFS_SERVER\u_new" on Vserver "vs1".
```

```
cluster1::> vserver cifs users-and-groups local-user rename -user-name CIFS_SERVER\u_old -new-user
-name CIFS_SERVER\u_new -vserver vs1
```

```
[[IDc3daaa1b7634a586c8d0e08eb08a4b0c]]
= vserver cifs users-and-groups local-user set-password
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set a password for a local user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-user set-password` command sets
the password for the specified local user. The password must meet the
following criteria:

* The password must be at least six characters in length.
* The password must not contain user account name.
* The password must contain characters from three of the following four
categories:

  * English uppercase characters (A through Z)
  * English lowercase characters (a through z)
  * Base 10 digits (0 through 9)
  * Special characters: ~, !, @, #, 0, ^, , *, _, -, +, =, `, \, |, (, ),
[, ], :, ;, ", ', <, >, ,, ., ?, /





== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-name <CIFS name>`` - User Name::
This specifies the user name.

== Examples

The following example sets the password for the local user
"CIFS_SERVER\u1" associated with Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups local-user set-password -user-name CIFS_SERVER\u1 -vserver vs1

Enter the new password:
Confirm the new password:

```
 + + The following example attempts to set the password but fails because
the password did not meet password complexity requirements.
```

cluster1::> vserver cifs users-and-groups local-user set-password -user-name CIFS_SERVER\u1 -vserver vs1

Enter the new password:
Confirm the new password:

Error: command failed: The password does not meet the password complexity
requirements. Refer to the man page for details.

```
[[ID45e4b18b6bb19e1b4f6ebba821314b06]]
= vserver cifs users-and-groups local-user show-membership
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local users' membership information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-user show-membership` command
displays the membership of local users.

== 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.

``[-vserver <vserver name>]`` - Vserver::
```

```
If this parameter is specified, the command displays local user membership
information for local users that are associated with the specified
Vserver.

``[-user-name <CIFS name>]`` - User Name::
If this parameter is specified, the command displays local user membership
information for a local user that matches the specified user name.

``[-membership <CIFS name>,...]`` - Local Group That This User is a Member
of::
If this parameter is specified, the command displays local user membership
information for the local group of which this local user is a member.

== Examples

The following example displays the membership information of all local
users; user "CIFS_SERVER\Administrator" is a member of
"BUILTIN\Administrators" group, and "CIFS_SERVER\u1" is a member of
"CIFS_SERVER\g1" group.
```

cluster1::> vserver cifs users-and-groups local-user show-membership
Vserver User Name Membership
---------- ---------- ----------
vs1 CIFS_SERVER\Administrator BUILTIN\Administrators
CIFS_SERVER\u1 CIFS_SERVER\g1
2 entries were displayed.

```
[[ID05b599cceb7a970f03f6eb86e7de614e]]
= vserver cifs users-and-groups local-user show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local users

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups local-user show` command displays local
users and their attributes.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays information only
about local users that match the specified Vserver name.

``[-user-name <CIFS name>]`` - User Name::
If this parameter is specified, the command displays information only
about local users that match the specified user name.

``[-full-name <TextNoCase>]`` - Full Name::
If this parameter is specified, the command displays information only
about local users that match the specified full name.

``[-description <TextNoCase>]`` - Description::
If this parameter is specified, the command displays information only
about local users that match the specified description.

``[-is-account-disabled {true|false}]`` - Is Account Disabled::
If this parameter is specified, the command displays information only
about local users that match the status specified.

== Examples

The following example displays information about all local users.
```

```
cluster1::> vserver cifs users-and-groups local-user show
Vserver User Name Full Name Description
------------ -------------------------- ------------------- -------------
vs1 CIFS_SERVER\Administrator James Raynor Built-in administrator account
vs1 CIFS_SERVER\u1 Sarah Kerrigan
2 entries were displayed.
```

```
[[IDcc94a135d48227819d5fa309251c1a2e]]
= vserver cifs users-and-groups privilege add-privilege
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add local privileges to a user or group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups privilege add-privilege` command adds
privileges to a local or Active Directory user or group.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-or-group-name <CIFS name>`` - User or Group Name::
This specifies the name of the local or Active Directory user or group.

``-privileges <Privilege>,...`` - Privileges::
This specifies the list of privileges to be associated with this user or
group.

== Examples

The following example adds the privileges "SeTcbPrivilege" and
"SeTakeOwnershipPrivilege" to the user "CIFS_SERVER\u1".
```

cluster1::> vserver cifs users-and-groups privilege add-privilege -vserver vs1 -user-or-group-name CIFS_SERVER\u1 -privileges SeTcbPrivilege,SeTakeOwnershipPrivilege

```
[[IDf9b478f9444e5a20f05634b18ee0205b]]
= vserver cifs users-and-groups privilege remove-privilege
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove privileges from a user or group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups privilege remove-privilege` command
removes privileges from a local or Active Directory user or group. This
command creates a new or modifies an existing privilege entry.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-or-group-name <CIFS name>`` - User or Group Name::
This specifies the name of the local or Active Directory user or group.

``-privileges <Privilege>,...`` - Privileges::
This specifies the list of privileges to be removed from this user or
group.

== Examples

The following example removes the previously added "SeTcbPrivilege" and
"SeTakeOwnershipPrivilege" privileges from the user "CIFS_SERVER\u1".
```

cluster1::> vserver cifs users-and-groups privilege show
Vserver User or Group Name Privileges
--- --- ---
vs1 CIFS_SERVER\u1 SeTcbPrivilege
SeTakeOwnershipPrivilege

cluster1::> vserver cifs users-and-groups privilege remove-privilege -vserver vs1 -user-or-group-name
CIFS_SERVER\u1 -privileges SeTcbPrivilege,SeTakeOwnershipPrivilege

cluster1::> vserver cifs users-and-groups privilege show
Vserver User or Group Name Privileges

---------- ---------- ----------
vs1 CIFS_SERVER\u1 -

```
 + + The following example removes "SeBackupPrivilege" from the group
"BUILTIN\Administrators".
```

cluster1::> vserver cifs users-and-groups privilege show
This table is currently empty.

cluster1::> vserver cifs users-and-groups privilege remove-privilege -vserver vs1 -user-or-group-name BUILTIN\Administrators -privileges SeBackupPrivilege

cluster1::> vserver cifs users-and-groups privilege show
Vserver User or Group Name Privileges
---------- ---------- ----------
vs1 BUILTIN\Administrators SeRestorePrivilege
SeSecurityPrivilege
SeTakeOwnershipPrivilege

```
[[ID381f40c6e819b220e750af7d4a73e565]]
= vserver cifs users-and-groups privilege reset-privilege
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset local privileges for a user or group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups privilege reset-privilege` command
resets privileges of a local or Active Directory user or group.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver.

``-user-or-group-name <CIFS name>`` - User or Group Name::
This specifies the name of the local or Active Directory user or group.

== Examples

The following example resets the privileges for the local user
"CIFS_SERVER\u1". This operation removes the privilege entry, if any,
associated with the local user "CIFS_SERVER\u1".
```

cluster1::> vserver cifs users-and-groups privilege show
Vserver User or Group Name Privileges
---------- ---------- ----------
vs1 CIFS_SERVER\u1 SeTakeOwnershipPrivilege
SeRestorePrivilege

cluster1::> vserver cifs users-and-groups privilege reset-privilege -vserver vs1 -user-or-group-name
CIFS_SERVER\u1

cluster1::> vserver cifs users-and-groups privilege show
This table is currently empty.

```
 + + The following example resets the privileges for the group
"BUILTIN\Administrators", effectively removing the privilege entry.
```

cluster1::> vserver cifs users-and-groups privilege show
Vserver User or Group Name Privileges
--- --- ---
vs1 BUILTIN\Administrators SeRestorePrivilege
SeSecurityPrivilege
SeTakeOwnershipPrivilege

cluster1::> vserver cifs users-and-groups privilege reset-privilege -vserver vs1 -user-or-group-name BUILTIN\Administrators

cluster1::> vserver cifs users-and-groups privilege show
This table is currently empty.

```
[[IDc6f665ba00c7fa0a16beefd45a44fdd4]]
= vserver cifs users-and-groups privilege show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display privileges

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver cifs users-and-groups privilege show` command displays
privilege overrides assigned to local or Active Directory users or groups.

== 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.

``[-vserver <vserver name>]`` - Vserver::
```

```
If this parameter is specified, the command displays information only
about privilege overrides assigned to local or Active Directory users or
groups that match the specified Vserver name.

``[-user-or-group-name <CIFS name>]`` - User or Group Name::
If this parameter is specified, the command displays information only
about privilege overrides assigned to local or Active Directory users or
groups that match the specified user name or group name.

``[-privileges <Privilege>,...]`` - Privileges::
If this parameter is specified, the command displays information only
about privilege overrides assigned to local or Active Directory users or
groups that match the specified privileges.

== Examples

The following example displays all privileges explicitly associated with
local or Active Directory users or groups for Vserver "vs1".
```

cluster1::> vserver cifs users-and-groups privilege show -vserver vs1
Vserver User or Group Name Privileges
--- --- ---
vs1 BUILTIN\Administrators SeTakeOwnershipPrivilege
SeRestorePrivilege

```
:leveloffset: -1


= vserver config-replication commands

:leveloffset: +1


[[IDec1f6faf5937d60091c36c63541f0290]]
= vserver config-replication pause
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Temporarily pause Vserver configuration replication

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Vserver domain locking functionality locks the Vserver while Vserver DM is
recording a configuration baseline. This command aborts the ongoing
baseline generation activity, unlocks the Vserver and temporarily pauses
configuration replication for the Vserver. Command confirmations have to
be enabled to execute this command. The time at which replication resumes
is displayed after successful completion of the command. Configuration
changes made after executing this command are not replicated to the
partner cluster. If a disaster occurs during this time, the configuration
changes made are lost. Replication can be manually resumed by executing
the xref:{relative_path}vserver-config-replication-resume.html[vserver
config-replication resume] command.

== Parameters

``-vserver <vserver name>`` - Vserver name [.advaced-priv]#(privilege:
advanced)#::


== Examples
```

cluster::> vserver config-replication pause -vserver vs1
Vserver configuration replication will be paused, then automatically resumed after five minutes.
Manually resume configuration replication by running the "vserver config-replication resume -vserver vs1"

command.
Do you want to continue ? {y|n}: y
Vserver configuration replication is paused and will be resumed at: 5/24/2014 06:11:23

```
== Related Links

* xref:{relative_path}vserver-config-replication-resume.html[vserver
config-replication resume]




[[ID62b528588aecf99c8bcb4cda55fd2903]]
= vserver config-replication resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume Vserver configuration replication

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command resumes configuration replication of the Vserver which was
temporarily paused by using the xref:{relative_path}vserver-config-
replication-pause.html[vserver config-replication pause] command.
Successful completion of the command ensures that configuration
replication has been resumed for the Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver name [.advaced-priv]#(privilege:
advanced)#::


== Examples
```

cluster::> vserver config-replication resume -vserver vs1

```
== Related Links

* xref:{relative_path}vserver-config-replication-pause.html[vserver
```

```
config-replication pause]




[[ID104561983ae5fe8e30d840c5d6fa2b53]]
= vserver config-replication show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Vserver configuration replication resume time

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver config-replication show` command displays the time at which
the configuration replication resumes for the Vserver.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays resume time for the
specified Vserver.

``[-resume-time <MM/DD/YYYY HH:MM:SS>]`` - Replication resume time
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays Vservers whose
configuration replications are resumed at the specified resume time.

== Examples
```

```
cluster::> vserver config-replication show
Replication
Vserver Resume Time
----------- ------------------
vs1 12/9/2014 03:18:48
```

```
:leveloffset: -1


= vserver consistency-group commands

:leveloffset: +1


[[IDc2afc011abeab30e4b0b3325e443719d]]
= vserver consistency-group attach
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Attach a consistency group to an existing parent consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group attach` command can be used to attach a
consistency group to a parent consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
that is to be attached.

``-consistency-group <text>`` - Consistency Group::
This parameter specifies the consistency group that is to be attached.

``-parent-consistency-group <text>`` - Parent Consistency Group::
This parameter specifies the parent consistency group to be attached to.

``[-new-name <text>]`` - New Name for the Consistency Group::
This parameter optionally specifies a new name for the attached
consistency group.

== Examples

The following command attaches the consistency group singleCG to a parent
consistency group parentCG in Vserver vs0.
```

cluster1::> vserver consistency-group attach -vserver vs0 -consistency-group childCG -parent-consistency -group parentCG
[Job 174] Job succeeded: Success

```
 + The following command attaches the consistency group singleCG to a
parent consistency group parentCG in Vserver vs0, which is renamed to
childCG.
```

cluster1::> vserver consistency-group attach -vserver vs0 -consistency-group childCG -parent-consistency -group parentCG -new-name childCG
(vserver consistency-group attach)
[Job 174] Job succeeded: Success

```
[[ID35b21f23424c4a20cccf6b16db38f015]]
= vserver consistency-group create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group create` command can be used to create a
consistency group using existing consistency groups or volumes, or by
creating new volumes.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver in which the consistency group is to
be created.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the name of the consistency group which is to be
created.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the name of the existing parent consistency group
in which the consistency group is to be created. If the parent consistency
group does not exist, it will be created.
```

```
``[-storage-service <text>]`` - Storage Service::
This parameter specifies the storage service name. If not specified, the
default value is the most performant for the platform.

``[-qos-policy <text>]`` - QoS Policy Group::
This parameter specifies the QoS policy to be applied to the consistency
group during creation.

``[-tiering-placement-rules <FabricPool Placement Preferences>]`` -
Tiering Placement Rules::
This parameter specifies the storage tiering placement rules for the
consistency group.

``[-tiering-policy <Tiering Policy>]`` - Tiering Policy::
This parameter specifies the tiering policy to be applied to the
consistency group during creation.

``[-object-stores <text>,...]`` - Object Store Name::
This parameter specifies the remote object stores to used for placement.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Copy Policy::
This parameter specifies the snapshot policy to be applied to the
consistency group during creation.

{ ``[-application-type <Application type for the parent or top level
CG>]`` - Application Type::
This parameter specifies the application type for the parent consistency
group.

| ``[-application-component-type <Application component type for child
CG>]`` - Application Component Type }::
This parameter specifies the application component type of the child
consistency group.

{ ``[-consistency-groups <text>,...]`` - Consistency Groups::
This parameter optionally specifies a comma separated list of existing
consistency groups under the Vserver.

| ``[-volumes <text>]`` - Volume Names::
This parameter specifies a filter to choose any existing volumes in the
Vserver to add to the new consistency group.

| ``[-volume-prefix <volume name>]`` - Volume Name Prefix::
This parameter specifies a volume prefix to be added to the volume name
for new volumes created in the new consistency group.
```

```
``[-volume-count <integer>]`` - Number of Volumes to Create }::
This parameter specifies the number of new volumes to be created in the
new consistency group.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Provisioned Size::
This parameter specifies the size of each new volume that is to be created
in the consistency group. If `-lun` or `-namespace` parameter is
specified, this refers to the size of each LUN or namespace.

{ ``[-luns <text>]`` - Existing LUN Names::
This parameter specifies a filter to choose any existing LUNs in the
Vserver to add to the new consistency group.

``[-lun-prefix <text>]`` - LUN Name Prefix::
This parameter specifies the name of the LUN to be created in the
consistency group. If the `-lun-count` parameter is specified this field
is treated as prefix.

``[-lun-count <integer>]`` - Number of LUNs to Create::
This parameter specifies the number of new LUNs to be created in the
consistency group.

``[-lun-os-type <LUN Operating System Format>]`` - LUN Operating System
Type::
This parameter specifies the OS type for the new LUNs.

``[-igroup <text>]`` - iGroup Name::
This parameter specifies the name of the initiator group.

| ``[-namespaces <text>]`` - Existing Namespace Names::
This parameter specifies a filter to choose any existing namespaces in the
Vserver to add to the new consistency group.

``[-namespace-prefix <text>]`` - Namespace Name Prefix::
This parameter specifies the name of the namespace to be created in the
consistency group. If the `-namespace-count` parameter is specified this
field is treated as prefix.

``[-namespace-count <integer>]`` - Number of Namespaces to Create::
This parameter specifies the number of new namespaces to be created in the
consistency group.

``[-namespace-os-type {aix|linux|vmware|windows}]`` - NVME Operating
System Type::
This parameter specifies the OS type for the new namespaces.

``[-subsystem <text>]`` - Subsystem Name::
```

```
This parameter specifies the name of the nvme subsystem.

| ``[-export-policy <text>]`` - Export Policy Name::
This parameter specifies the name of the export policy to be associated
with the newly created volumes.

``[-nas-gid <integer>]`` - NAS Group ID::
This parameter specifies the UNIX group ID of the newly created volumes.

``[-nas-path <text>]`` - Junction Path::
This parameter specifies the mount path for the newly created volumes.

``[-nas-junction-parent-volume <volume name>]`` - Junction Parent Volume
Name::
This parameter specifies the name of the parent volume that contains the
junction inode of this volume.

``[-nas-security-style <security style>]`` - NAS Security Style::
This parameter specifies the security style associated with the newly
created volumes.

``[-nas-uid <integer>]`` - NAS User ID::
This parameter specifies the UNIX user ID of the newly created volumes.

``[-nas-unix-permissions <unix perm>]`` - NAS UNIX Permissions::
This parameter specifies the UNIX permissions for the newly created
volumes.

``[-cifs-share <Share>]`` - Volume CIFS Share Name::
This parameter specifies the name of the CIFS share for each volume in the
new consistency group.

``[-cifs-share-acl-user-or-group <text>]`` - CIFS User/Group Name::
This parameter specifies the ACL user or group of the CIFS share for each
volume in the new consistency group.

``[-cifs-share-acl-win-unix-id <text>]`` - Windows SID or UNIX ID::
This parameter specifies the ACL windows or unix id of the CIFS share for
each volume in the new consistency group.

``[-cifs-share-acl-type {windows|unix-user|unix-group}]`` - CIFS User or
Group Type::
This parameter specifies the ACL type of the CIFS share for each volume in
the new consistency group.

``[-cifs-share-acl-permission <access rights>]`` - CIFS Access Type }::
This parameter specifies the ACL permission of the CIFS share for each
```

```
volume in the new consistency group.

== Examples

 + The following command creates a new parent consistency group parentcg
with existing consistency groups cg1 and cg2.
```

cluster1::> vserver consistency-group create -consistency-group parentcg -consistency-groups cg1,cg2
(vserver consistency-group create)
[Job 32] Job succeeded: Success

```
 + The following command creates a new consistency group singlecg with
existing volumes vol1 and vol2.
```

cluster1::> vserver consistency-group create -consistency-group singlecg -volumes vol1,vol2
(vserver consistency-group create)
[Job 33] Job succeeded: Success

```
 + The following command creates a new consistency group singlecg with two
new volumes each of size 1gb and volume name prefix with db_vols.
```

cluster1::> vserver consistency-group create -consistency-group singlecg -volume-prefix db_vols -volume
-count 2 -size 1gb
(vserver consistency-group create)
[Job 34] Job succeeded: Success

```
 + The following command creates a new consistency group singlecg2 with
two new volumes each of size 1gb and volume name prefix with db_vols and a
snapshot policy of default and application-type mongodb.
```

cluster1::> vserver consistency-group create -consistency-group singlecg -volume-prefix db_vols -volume
-count 2 -size 1gb -snapshot-policy default -application-type mongodb
(vserver consistency-group create)
[Job 35] Job succeeded: Success

```
 + The following command creates a new consistency group child1 under
existing parent consistency group parent1 and create two new volumes each
of size 1gb.
```

cluster1::> vserver consistency-group create -consistency-group child1 -parent-consistency-group parent1
-volume-count 2 -size 1gb
(vserver consistency-group create)
[Job 36] Job succeeded: Success

```
 + The following command creates a new consistency group child2 under
existing parent consistency group parent2 and creates two new volumes each
of size 1gb and volume name prefix with child2_volumes.
```

cluster1::> vserver consistency-group create -consistency-group child2 -parent-consistency-group parent2 -volume-prefix child2_volumes -volume-count 2 -size 1gb
(vserver consistency-group create)
[Job 37] Job succeeded: Success

```
 + The following command creates a new consistency group child3 under
existing parent consistency group parent2 and creates two new volumes vol1
and vol2 each of size 1gb.
```

cluster1::> vserver consistency-group create -consistency-group child3 -parent-consistency-group parent2 -volumes vol1,vol2 -size 1gb
(vserver consistency-group create)
[Job 38] Job succeeded: Success

```
 + The following command creates a new consistency group singlecg with two
new volumes each of size 1gb and volume name prefix with db_vols and CIFS
share share1 and ACL properties.
```

cluster1::> vserver consistency-group create -consistency-group singlecg -volume-prefix db_vols -volume -count 2 -size 1gb -cifs-share share1 -cifs-share-acl-type windows -cifs-share-acl-user-or-group Everyone -cifs -share-acl-permission Read -nas-path "/vol"
(vserver consistency-group create)
[Job 39] Job succeeded: Success

```
 + The following command creates a new consistency group cg1 with existing
LUNs lun1 and lun2.
```

cluster1::> vserver consistency-group create -consistency-group cg1 -luns lun1,lun2
(vserver consistency-group create)
[Job 33] Job succeeded: Success

```
 + The following command creates a new consistency group cg1 with existing
namespaces ns1 and ns2 using wildcards.
```

cluster1::> vserver consistency-group create -consistency-group cg1 -namespaces ns*
(vserver consistency-group create)
[Job 33] Job succeeded: Success

```
 + The following command creates a new consistency group cg1 with two new
LUNs each of size 1gb with OS type linux.
```

cluster1::> vserver consistency-group create -consistency-group cg1 -lun-prefix db_luns -lun-count 2 -lun-os -type linux -size 1gb
(vserver consistency-group create)
[Job 34] Job succeeded: Success

```
 + The following command creates a new consistency group cg1 with three
new namespaces each of size 1gb with OS type linux.
```

cluster1::> vserver consistency-group create -consistency-group cg1 -namespace-prefix db_ns -namespace -count 3 -namespace-os-type linux -size 1gb
(vserver consistency-group create)
[Job 34] Job succeeded: Success

```
[[IDce49a341f758fe652abaf98f6db636cd]]
= vserver consistency-group delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group delete` command can be used to delete a
consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
that is to be deleted.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group that is to be deleted.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group of the consistency
group that is to be deleted.

== Examples

The following command deletes the consistency group parentCG in Vserver
vs0.
```

cluster1::> vserver consistency-group delete -vserver vs0 -consistency-group parentCG -parent-consistency -group -
(vserver consistency-group delete)
Warning: Are you sure you want to delete consistency group "parentCG" in
Vserver "vs0" ? {y|n}: y
[Job 174] Job succeeded: Success
1 entry was deleted.

```
[[IDff2caa56537abd73070193f33b174457]]
= vserver consistency-group demote
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Demote a parent consistency group to become standalone consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group demote` command can be used to demote a
parent consistency group to be on its own, deleting its child consistency
groups.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
that is to be demoted.

``-parent-consistency-group <text>`` - Parent Consistency Group::
This parameter specifies the parent consistency group that is to be
demoted.

``[-new-name <text>]`` - New name for the Consistency Group::
This parameter optionally specifies a new name for the consistency group
after demotion.

== Examples

The following command demotes the consistency group parentCG in Vserver
vs0, which is renamed to singleCG at demotion.
```

cluster1::> vserver consistency-group demote -vserver vs0 -parent-consistency-group parentCG -new-name singleCG
(vserver consistency-group demote)
[Job 174] Job succeeded: Success

```
[[IDfeb0552cb1a4476c42617c29e180ab38]]
= vserver consistency-group detach
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Detach a child consistency group from an existing parent consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group detach` command can be used to detach a
child consistency group from its parent, to be on its own. If this was the
only child under that parent, the parent consistency group will be
deleted.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
that is to be detached.

``-parent-consistency-group <text>`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group.

``-consistency-group <text>`` - Consistency Group to Detach::
This parameter specifies the consistency group that is to be detached.

``[-new-name <text>]`` - New name for the Detached Consistency Group::
This parameter optionally specifies a new name for the detached
consistency group.

== Examples

The following command detaches the consistency group childCG from its
parent Consistency Group parentCG in Vserver vs0, which is renamed to
singleCG at removal.
```

cluster1::> vserver consistency-group detach -vserver vs0 -parent-consistency-group parentCG -consistency -group childCG -new-name singleCG
(vserver consistency-group detach)
[Job 174] Job succeeded: Success

```
[[IDcebaf05c95c0d8aa580bda7669a0c0ac]]
= vserver consistency-group modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the configuration of an existing consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group modify` command can be used to modify the
following attributes of a consistency group:

* Application component type
* Application type
* Snapshot policy
* Minimum SnapLock retention
* Maximum SnapLock retention
* Default SnapLock retention



== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
that is to be modified

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group that is to be modified.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Copy Policy::
This optionally specifies the snapshot policy for the consistency group.

{ ``[-application-type <Application type for the parent or top level
CG>]`` - Application Type::
This optionally specifies the application type for the parent consistency
group.
```

```
| ``[-application-component-type <Application component type for child
CG>]`` - Application Component Type }::
This optionally specifies the application component type for the child
consistency group.

== Examples

The following command modifies the snapshot policy of consistency group
childCg in parent consistency group parentCg in vserver vs0 to default
snapshot policy.
```

cluster1::> vserver consistency-group modify -vserver vs0 -consistency-group childCg -parent-consistency -group parentCg -snapshot-policy default
[Job 51] Job succeeded: Success
1 entry was modified.

```
[[IDbbcc33c619f926e9828f08820d50a8ae]]
= vserver consistency-group promote
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Promote a standalone consistency group to become parent consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group promote` command can be used to promote a
consistency group to a parent consistency group. A new child consistency
group will be created and associated with the newly promoted parent
consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
that is to be promoted.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group that is to be promoted.

``-child-consistency-group <text>`` - Child Consistency Group Name::
This parameter specifies the child consistency group which will get
created during promotion.

``[-new-name <text>]`` - New name for the Consistency Group::
This parameter optionally specifies a new name for the consistency group
after promotion.

== Examples

The following command promotes the consistency group singleCG in Vserver
vs0, which is renamed to parentCG at promotion gets assigned a new child
consistency group childCG at promotion.
```

cluster1::> vserver consistency-group promote -vserver vs0 -consistency-group singleCG -child-consistency -group childCG -new-name parentCG
(vserver consistency-group promote)

[Job 65] Job succeeded: Success

```
[[IDeb3283740addce71b42d8093d7244bc6]]
= vserver consistency-group show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of existing consistency groups

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The command displays information for consistency groups. Use the `
instance ` parameter to display additional consistency group details.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-statistics  ]``::
Selects the last recorded historic and real time statistics for
consistency group.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver Name::
Selects information about the consistency groups in the specified Vserver.

``[-consistency-group <text>]`` - Consistency Group Name::
Selects information about the specified consistency group.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
Selects information about the specified parent consistency group.

``[-qos-policy <text>]`` - QoS Policy Group::
Selects information about the consistency groups that have the specified
QoS policy.
```

```
``[-tiering-policy <Tiering Policy>]`` - Tiering Policy::
Selects information about the consistency groups that have the specified
tiering policy.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Copy Policy::
Selects information about the consistency groups that have the specified
snapshot policy.

``[-application-type <Application type for the parent or top level CG>]``
- Application Type::
Selects information about the consistency groups that have the specified
application type.

``[-application-component-type <Application component type for child
CG>]`` - Application Component Type::
Selects information about the consistency groups that have the specified
application component type.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Provisioned Size::
Selects information about the consistency groups that have the specified
size.

``[-state {online|creating|deleting|modifying|restoring}]`` - State::
Selects information about the consistency groups that have the specified
state.

``[-uuid <UUID>]`` - Consistency Group UUID::
Selects information about the consistency group that matches the specified
consistency group UUID.

``[-create-time <Date>]`` - Creation Time::
Selects information about the consistency groups that have the specified
creation time.

``[-space-available {<integer>[KB|MB|GB|TB|PB]}]`` - Available Space::
Selects information about the consistency groups that have the specified
available space.

``[-space-used {<integer>[KB|MB|GB|TB|PB]}]`` - Space Used::
Selects information about the consistency groups that have the specified
used space.

``[-replicated {true|false}]`` - Is Replicated::
Selects information about the consistency groups that have the specified
replicated status.

``[-replication-policy <text>]`` - Replication Policy Name::
```

```
Selects information about the consistency groups that have the specified
replication policy.

``[-replication-source {true|false}]`` - Is Replication Source::
Selects information about the consistency groups that have the specified
replication source.

``[-vdisk-type {luns|namespaces|mixed}]`` - Consistency Group Type::
Selects information about the consistency groups that have the specified
vdisk type.

``[-timestamp-metric <MM/DD/YYYY HH:MM:SS>]`` - Timestamp of the Sample::
Selects information about the last processed consistency group statistic
under the specified timestamp.

``[-duration <text>]`` - Duration of Time Used to Measure the Sample::
Selects information about the last processed consistency group statistic
under the specified duration.

``[-status-metric
{ok|error|partial_no_data|partial_no_response|partial_other_error|negative
_delta|not_found|backfilled_data|inconsistent_delta_time|partial_no_uuid|u
ndefined}]`` - Status of the Sample::
Selects information about the last processed consistency group statistic
under the specified status.

``[-available-space-metric <integer>]`` - Available Space in the
Consistency Group::
Selects information about the last processed consistency group statistic
under the specified space available.

``[-used-space-metric <integer>]`` - Space Used in the Consistency Group::
Selects information about the last processed consistency group statistic
under the specified space used.

``[-size-metric <integer>]`` - Size of the Consistency Group::
Selects information about the last processed consistency group statistic
under the specified size.

``[-latency-read-metric <integer>]`` - Latency of Read Operations (Us)::
Selects information about the last processed consistency group statistic
under the specified read latency.

``[-latency-write-metric <integer>]`` - Latency of Write Operations (Us)::
Selects information about the last processed consistency group statistic
under the specified write latency.
```

```
``[-latency-other-metric <integer>]`` - Latency of Other Operations (Us)::
Selects information about the last processed consistency group statistic
under the specified other latency.

``[-latency-total-metric <integer>]`` - Total Latency over All Operations
(Us)::
Selects information about the last processed consistency group statistic
under the specified total latency.

``[-iops-read-metric <integer>]`` - Number of Read Operations::
Selects information about the last processed consistency group statistic
under the specified read iops.

``[-iops-write-metric <integer>]`` - Number of Write Operations::
Selects information about the last processed consistency group statistic
under the specified write iops.

``[-iops-other-metric <integer>]`` - Number of Other Operations::
Selects information about the last processed consistency group statistic
under the specified other iops.

``[-iops-total-metric <integer>]`` - Total number of Operations::
Selects information about the last processed consistency group statistic
under the specified total iops.

``[-throughput-read-metric <integer>]`` - Throughput of Read Operations
(Bytes/s)::
Selects information about the last processed consistency group statistic
under the specified read throughput.

``[-throughput-write-metric <integer>]`` - Throughput of Write Operations
(Bytes/s)::
Selects information about the last processed consistency group statistic
under the specified write throughput.

``[-throughput-other-metric <integer>]`` - Throughput of Other Operations
(Bytes/s)::
Selects information about the last processed consistency group statistic
under the specified other throughput.

``[-throughput-total-metric <integer>]`` - Total Throughput over All
Operations (Bytes/s)::
Selects information about the last processed consistency group statistic
under the specified total throughput.

``[-timestamp-raw <MM/DD/YYYY HH:MM:SS>]`` - Timestamp of the Raw Sample::
Selects information about the real time consistency group statistic under
```

```
the specified timestamp.

``[-status-raw
{ok|error|partial_no_data|partial_no_response|partial_other_error|negative
_delta|not_found|backfilled_data|inconsistent_delta_time|partial_no_uuid|u
ndefined}]`` - Status of the Raw Sample::
Selects information about the real time consistency group statistic under
the specified status.

``[-available-space-raw <integer>]`` - Available Space in the Consistency
Group::
Selects information about the real time consistency group statistic under
the specified available space.

``[-used-space-raw <integer>]`` - Space Used in the Consistency Group::
Selects information about the real time consistency group statistic under
the specified used space.

``[-size-raw <integer>]`` - Size of the Consistency Group::
Selects information about the real time consistency group statistic under
the specified size.

``[-latency-read-raw <integer>]`` - Raw Latency over All Read Operations
(Us)::
Selects information about the real time consistency group statistic under
the specified read latency.

``[-latency-write-raw <integer>]`` - Raw Latency over All Write Operations
(Us)::
Selects information about the real time consistency group statistic under
the specified write latency.

``[-latency-other-raw <integer>]`` - Raw Latency of Other Operations
(Us)::
Selects information about the real time consistency group statistic under
the specified other latency.

``[-latency-total-raw <integer>]`` - Raw Total Latency over All Operations
(Us)::
Selects information about the real time consistency group statistic under
the specified total latency.

``[-iops-read-raw <integer>]`` - Number of Read Operations::
Selects information about the real time consistency group statistic under
the specified read iops.

``[-iops-write-raw <integer>]`` - Number of Write Operations::
```

```
Selects information about the real time consistency group statistic under
the specified write iops.

``[-iops-other-raw <integer>]`` - Number of Other Operations::
Selects information about the real time consistency group statistic under
the specified other iops.

``[-iops-total-raw <integer>]`` - Total number of Operations::
Selects information about the real time consistency group statistic under
the specified total iops.

``[-throughput-read-raw <integer>]`` - Number of Bytes of Read
Operations::
Selects information about the real time consistency group statistic under
the specified read throughput.

``[-throughput-write-raw <integer>]`` - Number of Bytes of Write
Operations::
Selects information about the real time consistency group statistic under
the specified write throughput.

``[-throughput-other-raw <integer>]`` - Number of Bytes of Other
Operations::
Selects information about the real time consistency group statistic under
the specified other throughput.

``[-throughput-total-raw <integer>]`` - Number of Total Bytes of All
Operations::
Selects information about the real time consistency group statistic under
the specified total throughput.

``[-clone-is-flexclone {true|false}]`` - Is Consistency Group a
FlexClone::
Selects information about the consistency group that have specified
FlexClone to be true.

``[-clone-source-consistency-group <text>]`` - Clone Source Consistency
Group::
Selects information about the consistency group that have specified clone
source consistency group.

``[-clone-source-consistency-group-uuid <UUID>]`` - Clone Source
Consistency Group UUID::
Selects information about the consistency group that have specified clone
source consistency group UUID.

``[-clone-source-snapshot <snapshot name>]`` - Snapshot Name of the Clone
```

```
Source::
Selects information about the consistency group that have specified clone
snapshot.

``[-clone-source-snapshot-uuid <UUID>]`` - Snapshot UUID of the Clone
Source::
Selects information about the consistency group that have specified clone
snapshot UUID.

``[-clone-parent-vserver <vserver name>]`` - Clone Parent Vserver Name::
Selects information about the cloned consistency groups with specified
clone parent Vserver

``[-clone-parent-vserver-uuid <UUID>]`` - Clone Parent Vserver UUID::
Selects information about the cloned consistency groups with specified
clone parent Vserver UUID

== Examples

The following command lists all the consistency groups on Vserver vs0.
```

vserver consistency-group show -vserver vs0
(vserver consistency-group show)
Parent
Consistency Consistency
Vserver Group Group State Size Available Used
---------- ------------- ------------- ------- ---------- ---------- ---------
svm1 cg1 - online 315.8MB 299.1MB 908KB
svm1 cg2 - online 105.3MB 99.72MB 288KB
svm1 cg3 cg1 online 315.8MB 299.1MB 908KB
3 entries were displayed.

The following command shows the statistics for consistency group cg1 on Vserver vs0.
SimpleClus::*> consistency-group show -vserver vs0 -consistency-group cg1 -statistic
(vserver consistency-group show)
Parent
Consistency Consistency
Vserver Group Group
---------- ------------- -------------
svm1 cg1 -
Last calculated statistic Value
------------------------- ------
timestamp-metric 11/13/2023 21:11:00
duration PT15S
status-metric ok
available-space-metric 156946432
used-space-metric 344064

size-metric 165568512
iops-other-metric 0
iops-read-metric 0
iops-write-metric 0
iops-total-metric 0
latency-other-metric 0
latency-read-metric 0
latency-write-metric 0
latency-total-metric 0
throughput-other-metric -
throughput-read-metric 0
throughput-write-metric 0
throughput-total-metric 0
Raw statistics Value
--- ---
timestamp-raw 11/13/2023 21:11:02
status-raw ok
available-space-raw 156946432
used-space-raw 344064
size-raw 165568512
iops-other-raw 0
iops-read-raw 0
iops-write-raw 0
iops-total-raw 0
latency-other-raw 0
latency-read-raw 0
latency-write-raw 0
latency-total-raw 0
throughput-other-raw -
throughput-read-raw 0
throughput-write-raw 0
throughput-total-raw 0

```
[[ID2d2581efa0d9d41d9edcbff4e8968a74]]
= vserver consistency-group clone create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a consistency group clone

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group clone create` command creates a clone of a
specified consistency group. It also takes in an optional parameter
source-snapshot-name to use for creating the clone. Only parent
```

```
consistency groups support cloning.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
whose clone is to be created.

``-clone-consistency-group <text>`` - Consistency Group Name of the
Clone::
This parameter specifies the name of the clone consistency group.

``-source-parent-consistency-group <text>`` - Parent Consistency Group
Name of the Source::
This parameter specifies the name of the consistency group to clone from
or its parent consistency if cloning a child consistency group.

``[-source-snapshot-name <snapshot name>]`` - Snapshot of the Source
Consistency Group::
This parameter optionally specifies the snapshot of the source parent
consistency group to be used for creating the clone.

``[-space-guarantee {none|volume}]`` - Space Guarantee Style::
This parameter optionally specifies the space guarantee style for the
FlexClone volumes in the clone consistency group. A value of ``_volume_``
reserves space on the aggregate for the entire volume. A value of
``_none_`` reserves no space on the aggregate, meaning that writes can
fail if the aggregate runs out of space. The default setting is inherited
from the parent volume in the source parent consistency group.

``[-split-after-clone <true>]`` - Split clone Volumes::
This parameter optionally specifies if the FlexClone volumes within the
clone consistency group would be split from their parent volume in the
source parent consistency group after the FlexClone volume is created.

``[-clone-volume-prefix <text>]`` - Clone Volume Name Prefix::
This parameter specifies an optional volume name prefix for cloned volumes
in the clone consistency group.

``[-clone-volume-suffix <text>]`` - Clone Volume Name Suffix::
This parameter specifies an optional volume name suffix for cloned volumes
in the clone consistency group.

== Examples

The following example creates a clone clone1 of the source parent
consistency group container1 on Vserver vs0. The space guarantee of the
```

```
cloned volumes under the clone consistency group is volume and the volume
names have a prefix clone1 and suffix of clone1end. The cloned volumes
have split initiated as True to split the clones from the parent volumes.
```

cluster1::> vserver consistency-group clone create -vserver vs0 -clone-consistency-group clone1 -source -parent-consistency-group container1 -space-guarantee volume -clone-volume-prefix clone1 -clone-volume -suffix clone1end -split-after-clone true -snaplock-type enterprise
(vserver consistency-group clone create)
[Job 264] Job succeeded: Success

```
 + The following example creates a clone clone_child_1 of the source child
consistency group child_1 in parent consistency group parentcg on Vserver
vs0.
```

cluster1::> vserver consistency-group clone create -vserver vs0 -clone-child-consistency-group clone_child_1 -source-parent-consistency-group parentcg
(vserver consistency-group clone create)
[Job 264] Job succeeded: Success

```
The following example creates a clone clone1 of the source parent
consistency group container1 on Vserver vs0. The storage unit names have a
prefix pre_clone and suffix of clone_end. The cloned consistency group
will have SnapLock type to be enterprise.
```

cluster1::> vserver consistency-group clone create -vserver vs0 -clone-consistency-group clone1 -source -parent-consistency-group container1 -clone-storage-unit-prefix pre_clone -clone-storage-unit-suffix clone_end -snaplock-type enterprise
(vserver consistency-group clone create)
[Job 264] Job succeeded: Success

```
[[ID6a5fd1dadaa9344769a841f87046d010]]
= vserver consistency-group lun show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of existing consistency group Luns

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The command displays information for consistency group LUNs. Use the `
instance ` parameter to display additional consistency group LUN details.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver Name::
Selects information about the consistency group LUNs in the specified
Vserver.

``[-consistency-group <text>]`` - Consistency Group::
Selects information about the consistency group LUNs in the specified
consistency group.

``[-parent-consistency-group <text>]`` - Parent Consistency Group::
Selects information about the consistency group LUNs in the specified
parent consistency group.

{ ``[-path <path>]`` - LUN Path::
Selects information about the consistency group LUN that matches the
specified LUN path.

| ``[-lun <text>]`` - LUN Name::
Selects information about the consistency group LUN that matches the
specified LUN name.

``[-volume <volume name>]`` - Volume Name::
Selects information about the consistency group LUNs that matches the
specified volume name.

``[-qtree <qtree name>]`` - Qtree Name }::
Selects information about the consistency group LUNs that matches the
specified qtree name.

``[-uuid <UUID>]`` - LUN UUID::
Selects information about the consistency group LUN that matches the
specified LUN UUID.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
Selects information about the consistency group LUNs that matches the
specified Vserver UUID.
```

```
``[-consistency-group-uuid <UUID>]`` - Consistency Group UUID::
Selects information about the consistency group LUNs that matches the
specified consistency group UUID.

``[-parent-consistency-group-uuid <UUID>]`` - Parent Consistency Group
UUID::
Selects information about the consistency group LUNs that matches the
specified parent consistency group UUID.

== Examples

The following command lists all the LUNS that are associated with a
consistency group.
```

cluster1::> consistency-group lun show
(vserver consistency-group lun show)
Parent
Consistency Consistency LUN
Vserver Group Group Path
---------- ---------- ---------- ----------
vs0 ChildCG_1 ParentCG /vol/ParentCG_01_vol_1/ChildCG_1_lun_1_1
vs0 ChildCG_2 ParentCG /vol/ParentCG_02_vol_1/ChildCG_2_lun_1_1
vs0 singleCG - /vol/singleCG_vol_1/singleCG_lun_1_1
vs0 singleCG - /vol/singleCG_vol_1/singleCG_lun_1_2
vs1 singleCG - /vol/singleCG_vol_1/singleCG_lun_1_1
vs1 singleCG - /vol/singleCG_vol_1/singleCG_lun_1_2
6 entries were displayed.

```
[[ID2936a81e5f8713b9e8a8e2ef57a3263f]]
= vserver consistency-group metrics show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of existing consistency group metrics

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays metrics for consistency groups.

== Parameters
```

```
{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
This specifies the additional consistency group metrics details.

``[-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Timestamp of the Sample::
Selects information about the consistency group metrics under the
specified timestamp.

``-consistency-group <text>`` - Consistency Group Name::
Selects information about the metrics of a specified consistency group.

``[-vserver <Vserver Name>]`` - Vserver Name::
Selects information about the consistency groups metrics in the specified
Vserver.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
Selects information about the metrics of a specified parent consistency
group.

``[-interval {1h|1d|1w|1m|1y}]`` - Interval of Consistency Group Metrics::
Selects information about the consistency group metrics under the
specified interval.

``[-status
{ok|error|partial_no_data|partial_no_response|partial_other_error|negative
_delta|not_found|backfilled_data|inconsistent_delta_time|partial_no_uuid|u
ndefined}]`` - Status of the Sample::
Shows information about the consistency group metrics under the specified
status.

``[-duration <text>]`` - Duration of Time Used to Measure the Sample::
Shows information about the consistency group metrics under the specified
duration.

``[-available-space <integer>]`` - Available Space in the Consistency
Group::
Shows information about the consistency group metrics under the specified
space availability.

``[-used-space <integer>]`` - Space Used in the Consistency Group::
Shows information about the consistency group metrics under the specified
used space.

``[-size <integer>]`` - Size of the Consistency Group::
```

```
Shows information about the consistency group metrics under the specified
size.

``[-latency-read <integer>]`` - Latency of Read Operations (us)::
Shows information about the consistency group metrics under the specified
read latency.

``[-latency-write <integer>]`` - Latency of Write Operations (us)::
Shows information about the consistency group metrics under the specified
write latency.

``[-latency-other <integer>]`` - Latency of Other Operations (us)::
Shows information about the consistency group metrics under the specified
other latency.

``[-latency-total <integer>]`` - Total Latency over All Operations (us)::
Shows information about the consistency group metrics under the specified
total latency.

``[-iops-read <integer>]`` - Number of Read Operations::
Shows information about the consistency group metrics under the specified
read iops.

``[-iops-write <integer>]`` - Number of Write Operations::
Shows information about the consistency group metrics under the specified
write iops.

``[-iops-other <integer>]`` - Number of Other Operations::
Shows information about the consistency group metrics under the specified
other iops.

``[-iops-total <integer>]`` - Total number of Operations::
Shows information about the consistency group metrics under the specified
total iops.

``[-throughput-read <integer>]`` - Throughput of Read Operations
(bytes/s)::
Shows information about the consistency group metrics under the specified
read throughput.

``[-throughput-write <integer>]`` - Throughput of Write Operations
(bytes/s)::
Shows information about the consistency group metrics under the specified
write throughput.

``[-throughput-other <integer>]`` - Throughput of Other Operations
(bytes/s)::
```

```
Shows information about the consistency group metrics under the specified
other throughput.

``[-throughput-total <integer>]`` - Total throughput over all Operations
(bytes/s)::
Shows information about the consistency group metrics under the specified
total throughput.

== Examples
```

The following command lists all the historical metrics for 1 hour interval for consistency group cg1 on Vserver vs0.
SimpleClus::*> vserver consistency-group metrics show -vserver vs0 -consistency-group cg1 -interval -1h
(vserver consistency-group metrics show)
Iops Latency Throughput
Timestamp Status Duration Total Total Total
--- --- --- --- --- ---
10/31/2023 19:33:45 ok PT15S 0 0 0
10/31/2023 19:33:30 partial_no_data PT15S 0 0 0
10/31/2023 19:33:15 partial_no_data PT15S 0 0 0
10/31/2023 19:33:00 partial_no_data PT15S 0 0 0

The following command lists all the timestamps in historical metrics iops for 1 day interval for consistency group cg1 on Vserver vs0 where write iops were greater than 100.
SimpleClus::*> vserver consistency-group metrics show -vserver vs0 -consistency-group cg1 -fields iops-write -interval 1d -iops-write >100
(vserver consistency-group metrics show)
timestamp iops-write
--- ---
"4/25/2024 15:35:00" 101
"4/24/2024 17:30:00" 105
"4/24/2024 17:25:00" 102
3 entries were displayed.

The following command lists all the timestamps in historical metrics for 1 day interval for consistency group cg1 on Vserver vs0 where other iops were between a given interval.
SimpleClus::*> vserver consistency-group metrics show -vserver vs0 -consistency-group cg1 -interval 1d -fields iops-other -iops-other 1078..1070|590..600
(vserver consistency-group metrics show)
timestamp iops-other
--- ---
"4/25/2024 13:55:00" 1078
"4/25/2024 13:50:00" 1074
"4/25/2024 13:45:00" 1070
"4/25/2024 03:55:00" 598
"4/25/2024 03:50:00" 594
5 entries were displayed.

The following command lists all the historical metrics for 1 day interval for consistency group cg1 on Vserver vs0 with a given timestamp when throughput write was more than 1102.
SimpleClus::*> vserver consistency-group metrics show -vserver vs0 -consistency-group cg1-interval 1d -fields throughput-total -timestamp "4/25/2024 14:30:00".."4/25/2024 13:25:00" -throughput-write >=1102
(vserver consistency-group metrics show)
timestamp throughput-total
--- ---
"4/25/2024 14:30:00" 1106
"4/25/2024 14:25:00" 1102
2 entries were displayed.

```
[[ID416b67b880346e44f36de552816ff8e1]]
= vserver consistency-group namespace show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of existing consistency group namespaces

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The command displays information for consistency group namespaces. Use the
` instance ` parameter to display additional consistency group namespace
details.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver Name::
Selects information about the consistency group namespaces in the
specified Vserver.

``[-consistency-group <text>]`` - Consistency Group Name::
Selects information about the consistency group namespaces in the
```

```
specified consistency group.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
Selects information about the consistency group namespaces in the
specified parent consistency group.

{ ``[-path <path>]`` - Namespace Path::
Selects information about the consistency group namespace that matches the
specified namespace path.

| ``[-namespace <text>]`` - Namespace Name::
Selects information about the consistency group namespace that matches the
specified namespace.

``[-volume <volume name>]`` - Volume Name::
Selects information about the consistency group namespaces that matches
the specified volume name.

``[-qtree <qtree name>]`` - Qtree Name }::
Selects information about the consistency group namespaces that matches
the specified qtree name.

``[-uuid <UUID>]`` - Namespace UUID::
Selects information about the consistency group namespace that matches the
specified namespace UUID.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
Selects information about the consistency group namespaces that matches
the specified Vserver UUID.

``[-consistency-group-uuid <UUID>]`` - Consistency GroupUUID::
Selects information about the consistency group namespaces that matches
the specified consistency group UUID.

``[-parent-consistency-group-uuid <UUID>]`` - Parent Consistency Group
UUID::
Selects information about the consistency group namespaces that matches
the specified parent consistency group UUID.

== Examples

The following command lists all the namespaces that are associated with a
consistency group.
```

cluster1::> sti115nscluster-1::*> consistency-group namespace show
(vserver consistency-group namespace show)

```
Parent
Consistency Consistency Namespace
Vserver Group Group Path
---------- ------------- ------------- -------------
vs0 cg_test - /vol/vol_test/qtree_test/ns_test
vs0 child1 parent_nvme /vol/newVolnvme1/ns1_1
vs0 child1 parent_nvme /vol/newVolnvme1/ns1_2
vs0 child2 parent_nvme /vol/newVolnvme2/ns2_1
vs0 child2 parent_nvme /vol/newVolnvme2/ns2_2
vs0 single_nvme - /vol/single_nvme_1/ns1_1
vs0 single_nvme - /vol/single_nvme_1/ns1_2
vs1 child1 parent_nvme /vol/newVolnvme1/ns1_1
vs1 child1 parent_nvme /vol/newVolnvme1/ns1_2
vs1 child2 parent_nvme /vol/newVolnvme2/ns2_1
vs1 child2 parent_nvme /vol/newVolnvme2/ns2_2
11 entries were displayed.
```

```
[[ID7f078b073874f7e2ac096ed6f95a07bb]]
= vserver consistency-group snapshot commit
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Commit a 2 phase Snapshot copy for a consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group snapshot commit` command commits a 2-phase
snapshot of a specified consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
where the snapshot is to be committed.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group where a snapshot is to be
committed.

``-snapshot <snapshot name>`` - Snapshot Copy Name::
This parameter specifies the name of the snapshot that is to be committed.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group where a snapshot is
to be committed.

== Examples

The following example commits a 2-phase snapshot named snap1 on a child
consistency group named cg1 in parent consistency group parentCg on a
Vserver named vs0.
```

cluster1::> vserver consistency-group snapshot commit -vserver vs0 -consistency-group cg1 -parent -consistency-group parentCg -snapshot snap1
(vserver consistency-group snapshot commit)
[Job 100] Job succeeded: Success

```
[[ID9133acd100ecd37640c51cb8d2d634e2]]
= vserver consistency-group snapshot create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new consistency group Snapshot copy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group snapshot create` command creates a snapshot
of a specified consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
where the snapshot is to be created.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group where a snapshot is to be
created.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group where a snapshot is
to be created.

``-snapshot <snapshot name>`` - Snapshot Copy Name::
This parameter specifies the name of the snapshot that is to be created.

``[-consistency-type {crash|application}]`` - Consistency Type::
This parameter specifies the consistency level of a snapshot to be
created. The default value is ``_crash_`` .

``[-comment <text>]`` - Comment::
This parameter specifies the comment associated with a snapshot to be
created.

``[-snapmirror-label <text>]`` - Snapmirror Label::
This parameter specifies the label associated with a snapshot to be
created.
```

```
``[-write-fence {true|false}]`` - Take write fence::
This parameter specifies if a write fence is taken on the volumes in the
consistency group while creating a snapshot.

== Examples

The following example creates a snapshot named snap1 on a child
consistency group named cg1 in parent consistency group parentCg on a
Vserver named vs0. The snapshot has a comment "Feb golden image", a
SnapMirror label "Label" and is crash-consistent.
```

cluster1::> vserver consistency-group snapshot create -vserver vs0 -consistency-group cg1 -parent -consistency-group parentCg -snapshot snap1 -comment "Feb golden image" -snapmirror-label "Label" -consistency-type crash
(vserver consistency-group snapshot create)
[Job 100] Job succeeded: Success

```
[[ID262f2850915f6292dc72c8b81f3fe17f]]
= vserver consistency-group snapshot delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an existing consistency group Snapshot Copy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group snapshot delete` command deletes a snapshot
of a specified consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
on which the snapshot is to be deleted.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group where a snapshot is to be
deleted.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group where a snapshot is
to be deleted.

``-snapshot <snapshot name>`` - Snapshot Copy Name::
This parameter specifies the name of the snapshot that is to be deleted.

== Examples

The following example deletes a snapshot named snap1 on a child
consistency group named cg1 in parent consistency group parentCg on a
Vserver named vs0.
```

cluster1::> vserver consistency-group snapshot delete -vserver vs0 -consistency-group cg1 -parent -consistency-group parentCg -snapshot snap1
(vserver consistency-group snapshot delete)
Warning: Deleting a snapshot permanently removes data that is stored only in that snapshot. Are you sure you want to delete snapshot

"snap1" for consistency group "cg1" in Vserver "vs0" ? {y|n}: y

```
[[IDfc494592be604cb98a90fef8c4787eb4]]
= vserver consistency-group snapshot modify-snaplock-expiry-time
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify SnapLock expiry time of a consistency group snapshot

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group snapshot modify-snaplock-expiry-time`
command modifies a snapshot of a specified consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
on which the snapshot is to be modified.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group where a snapshot is to be
modified.

``-snapshot <snapshot name>`` - Snapshot Name::
This parameter specifies the name of the snapshot that is to be modified.

``-expiry-time {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm] | infinite}`` - SnapLock
Expiry Time::
Specifies the new SnapLock expiry that is applied to snapshot locked by
SnapLock.
+
If this option value is ``_infinite_`` , then a retention period that
never expires is applied to the snapshot.

== Examples

The following example extends the retention period of a snapshot
``_snap1_`` to ``_"03/03/2020 00:00:00"_`` :
```

cluster1::> vserver consistency-group snapshot modify-snaplock-expiry-time -vserver vs1 -consistency-group dst -snapshot s1 -expiry-time "03/03/2020 00:00:00"
cluster1::>

```
The following example extends the retention period of a snapshot
``_snap2_`` to ``_infinite_`` :
```

cluster1::> vserver consistency-group snapshot modify-snaplock-expiry-time -vserver vs1 -consistency-group dst -snapshot s2 -expiry-time infinite
cluster1::>

```

```

cluster1::> vserver consistency-group snapshot show -vserver vs1 -fields snaplock-expiry-time
(vserver consistency-group snapshot show)
vserver consistency-group parent-consistency-group snapshot snaplock-expiry-time
------- ---------- ---------- -------- ----------
vs1 dst - s1 3/3/2020 00:00:00 +05:30
vs1 dst - s2 infinite

```
[[ID9ab112fe78a414bd0649df35b1d030c1]]
= vserver consistency-group snapshot restore
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore a consistency group to a specified Snapshot copy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group snapshot restore` command restores a
snapshot of a specified consistency group. This replaces the current
working copy of the volume in the consistency group with the snapshot that
results in a loss of all changes made since the snapshot was created.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This specifies the Vserver that contains the consistency group on which
the specified snapshot to be restored is saved.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group where a snapshot is to be
restored.

``-snapshot <snapshot name>`` - Snapshot Copy Name::
This parameter specifies the name of the snapshot that is to be restored.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group where a snapshot is
to be restored.

== Examples

The following example restores a snapshot named snap1 on a consistency
group named childCG on parent consistency group parentCG and is located on
a Vserver named vs0.
```

cluster1::> vserver consistency-group snapshot restore -vserver vs0 -consistency-group childCG -parent -consistency-group parentCG -snapshot snap1
(vserver consistency-group snapshot restore)

[Job 100] Job succeeded: Success

```
[[ID60ffeb84c201712a5cfc379b2ce54446]]
= vserver consistency-group snapshot show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of existing consistency group Snapshot Copies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The command displays information for consistency group snapshots. Use the
` instance ` parameter to display additional consistency group snapshot
details.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that will be displayed.

| ``[-instance  ]`` }::
Selects detailed information about all fields about the consistency group
snapshots.

``[-vserver <Vserver Name>]`` - Vserver Name::
Selects information about the consistency group snapshots with a specified
Vserver.

``[-consistency-group <text>]`` - Consistency Group Name::
Selects information about the consistency group snapshots with a specified
consistency group.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
Selects information about the consistency group snapshots with a specified
parent consistency group.

``[-snapshot <snapshot name>]`` - Snapshot Copy Name::
Selects information about the specified consistency group snapshot.

``[-consistency-type {crash|application}]`` - Consistency Type::
```

```
Selects information about the consistency group snapshots with a specified
consistency type.

``[-comment <text>]`` - Comment::
Selects information about the consistency group snapshots with a specified
comment.

``[-snapmirror-label <text>]`` - Snapmirror Label::
Selects information about the consistency group snapshots with a specified
SnapMirror label.

``[-create-time <Date>]`` - Create Time::
Selects information about the consistency group snapshots with a specified
create time.

``[-pg-generation <integer>]`` - Protection Group Generation::
Selects information about the consistency group snapshots with a specified
pg generation.

``[-is-partial {true|false}]`` - Is Snapshot Copy Partial?::
Selects information about the consistency group snapshots with a specified
partial state.

``[-snapshot-volumes <volume name>,...]`` - List of Volume Names with this
Snapshot Copy::
Selects information about the consistency group snapshots with specified
snapshot volumes.

``[-missing-volumes <volume name>,...]`` - List of Volume Names Missing
this Snapshot Copy::
Selects information about the consistency group snapshots with specified
missing volumes.

``[-snapshot-uuid <UUID>]`` - Snapshot UUID::
Selects information about the consistency group snapshot that matches the
specified snapshot UUID.

``[-write-fence {true|false}]`` - Take write fence::
Selects information about the consistency group snapshots that matches the
specified write fence value.

== Examples

The following command lists all the snapshots that are associated with
consistency groups on Vserver svm1.
```

```
cluster1::> vserver consistency-group snapshot show -vserver svm1
(vserver consistency-group snapshot show)
Parent
Consistency Consistency Create
Vserver Group Group Snapshot Time
---------- ------------- ------------- ------------- ---------------
svm1 cg1 - snap1 Thu Jun 08 12:00:00 2023
svm1 cg2 - snap2 Thu Jun 08 1:00:00 2023
2 entries were displayed.
```

```
[[ID173e4dc17f0572e428ff6c27421bcb09]]
= vserver consistency-group snapshot start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start a 2 phase Snapshot copy for a consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group snapshot start` command starts a 2-phase
snapshot of a specified consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
where the snapshot is to be started.

``-consistency-group <text>`` - Consistency Group Name::
This parameter specifies the consistency group where a snapshot is to be
started.

``-snapshot <snapshot name>`` - Snapshot Copy Name::
This parameter specifies the name of the snapshot that is to be started.

``[-parent-consistency-group <text>]`` - Parent Consistency Group Name::
This parameter specifies the parent consistency group where a snapshot is
to be started for a child consistency group.

``[-consistency-type {crash|application}]`` - Consistency Type::
This parameter optionally specifies the consistency level of a snapshot.
```

```
The default value is ``_crash_`` .

``[-comment <text>]`` - Comment::
This parameter optionally specifies the comment associated with a
snapshot.

``[-snapmirror-label <text>]`` - Snapmirror Label::
This parameter optionally specifies the label associated with a snapshot.

``[-start-timeout <integer>]`` - Timeout for 2-phase snapshot::
This parameter optionally specifies the timeout limit for the snapshot to
be started. The default value is ``_7_`` seconds.

``[-write-fence {true|false}]`` - Take write fence::
This parameter specifies if a write fence is taken on the volumes in the
consistency group while creating a snapshot.

== Examples

 + The following example starts a 2-phase snapshot named snap2 on a child
consistency group named childCg in parent consistency group parentCg on a
Vserver named vs0. The snapshot has start-timeout set to "90".
```

cluster1::> vserver consistency-group snapshot start -vserver vs0 -consistency-group childCg -parent -consistency-group parentCg -snapshot snap2 -start-timeout 90
(vserver consistency-group snapshot start)

```
[[ID1d7cd4c74bbf4f968dcb93807a962138]]
= vserver consistency-group volume add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a volume to the consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group volume add` adds existing volumes to a
consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
to which the volumes will be added.

``-consistency-group <text>`` - Consistency Group::
This parameter specifies the consistency group to which the volumes will
be added.

``[-parent-consistency-group <text>]`` - Parent Consistency Group::
This parameter specifies the parent consistency group to which the volumes
will be added.

``-volume <text>`` - Volume::
This parameter specifies the volume to be added to the consistency group.

== Examples

The following example will add volumes vol1,vol2 to a child consistency
group named cg in parent consistency group parentCg on a Vserver named
vs0.
```

cluster1::> vserver consistency-group volume add -vserver vs0 -consistency-group cg -parent-consistency -group parentCg -volume vol1,vol2
(vserver consistency-group volume add)
[Job 100] Job succeeded: Success

```
[[ID95b60c9a30a5c474423ffda8e1e417e8]]
= vserver consistency-group volume create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new volume in a consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group volume create` command creates new volumes
in a consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
in which the volumes will be created.

``-consistency-group <text>`` - Consistency Group::
This parameter specifies the consistency group in which the volumes will
be created.

``[-parent-consistency-group <text>]`` - Parent Consistency Group::
This parameter specifies the parent of the consistency group in which the
volumes will be created.

``-volume <volume name>`` - Volume::
This parameter specifies the name of the volume to be created in the
consistency group. If the `-volume-count` parameter is specified this
field is treated as prefix.

``-size {<integer>[KB|MB|GB|TB|PB]}`` - Size::
This parameter specifies the size of each new volume that is to be created
in the consistency group. If `-lun` or `-namespace` parameter is
specified, this refers to the size of each LUN or namespace.

``[-volume-count <integer>]`` - Number of volumes::
This parameter specifies the number of new volumes to be created in the
consistency group.

{ ``[-export-policy <export policy name>]`` - Volume Export Policy::
```

```
This parameter specifies the name of the export policy to be associated
with the newly created volumes.

``[-nas-path <junction path>]`` - Junction Path::
This parameter specifies the junction path for mounting the volumes.

``[-nas-gid <integer>]`` - NAS Group ID::
This parameter specifies the NAS gid.

``[-nas-junction-parent-volume <volume name>]`` - Junction Parent Volume
Name::
This parameter specifies the NAS junction parent volume.

``[-nas-security-style <security style>]`` - NAS Security Style::
This parameter specifies the NAS security style.

``[-nas-uid <integer>]`` - NAS User ID::
This parameter specifies the NAS UID.

``[-nas-unix-permissions <unix perm>]`` - NAS UNIX Permissions::
This parameter specifies the NAS UNIX permissions.

``[-cifs-shares <Share>,...]`` - Volume CIFS Share Names::
This parameter specifies the name of the CIFS share to be created.

``[-cifs-share-acl-user-or-group <text>]`` - CIFS User/Group Name::
This parameter specifies the acl user or group of the CIFS share for each
volume in the new consistency group.

``[-cifs-share-acl-win-unix-id <text>]`` - Windows SID or UNIXID::
This parameter specifies the acl windows or unix id of the CIFS share for
each volume in the new consistency group.

``[-cifs-share-acl-type {windows|unix-user|unix-group}]`` - CIFS User or
Group Type::
This parameter specifies the acl type of the CIFS share for each volume in
the new consistency group.

``[-cifs-share-acl-permission <access rights>]`` - CIFS Access Type::
This parameter specifies the acl permission of the CIFS share for each
volume in the new consistency group.

| ``[-lun <text>]`` - LUN Name::
This parameter specifies the name of the LUN to be created in the
consistency group. If the `-lun-count` parameter is specified this field
is treated as prefix.
```

```
``[-lun-count <integer>]`` - Number of LUNs::
This parameter specifies the number of new LUNs to be created in the
consistency group.

``[-lun-os-type <LUN Operating System Format>]`` - OS Type::
This parameter specifies the OS type for the new LUNs.

``[-igroup <text>]`` - Igroup Name::
This parameter specifies the name of the initiator group.

| ``[-namespace <text>]`` - Namespace Name::
This parameter specifies the name of the namespace to be created in the
consistency group. If the `-namespace-count` parameter is specified this
field is treated as prefix.

``[-namespace-count <integer>]`` - Number of Namespaces::
This parameter specifies the number of new namespaces to be created in the
consistency group.

``[-namespace-os-type {aix|linux|vmware|windows}]`` - OS Type::
This parameter specifies the OS type for the new namespaces.

``[-subsystem <text>]`` - Subsytem Name }::
This parameter specifies the name of the nvme subsystem.

== Examples

 + The following command creates 2 new volumes in consistency group
singleCG.
```

cluster1::> vserver consistency-group volume create -vserver vs0 -consistency-group singleCG -parent -consistency-group - -volume vol_singleCG -size 20M -volume-count 2
(vserver consistency-group volume create)
[Job 100] Job succeeded: Success

```
 + The following command creates 2 new volumes in consistency group
singleCG with 2 new LUNs in each volume.
```

cluster1::> vserver consistency-group volume create -vserver vs0 -consistency-group singleCG -parent -consistency-group - -volume vol -size 20M -volume-count 2 -lun lun -lun-count 2 -lun-os-type linux -igroup ig1
(vserver consistency-group volume create)
[Job 101] Job succeeded: Success

```
 + The following command creates 2 new volumes in consistency group
singleCG with 2 new nvme namespaces in each volume.
```

cluster1::> vserver consistency-group volume create -vserver vs0 -consistency-group singleCG -parent -consistency-group - -volume vol -size 20M -volume-count 2 -namespace ns -namespace-count 2 -namespace -os-type linux -subsystem ss1
(vserver consistency-group volume create)
[Job 102] Job succeeded: Success

```
 + The following command creates a new volume in consistency group
singleCG with CIFS share along with acl properties.
```

cluster1::> vserver consistency-group volume create -vserver vs0 -consistency-group singleCG -parent -consistency-group - -volume vol_singleCG -size 20M -volume-count 2 -cifs-shares share1 -cifs-share-acl-type windows -cifs-share-acl-user-or-group Everyone -cifs-share-acl-permission Read -nas-path "/vol"
(vserver consistency-group volume create)
[Job 103] Job succeeded: Success

```
[[ID0e4f2d72e899b4d46ed64126bb36a1df]]
= vserver consistency-group volume reassign
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reasign a volume to a different consistency group.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group volume reassign` will reassign volumes from
one child consistency group to another child consistency group within a
parent consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
from which the volumes are reassigned.

``-consistency-group <text>`` - Consistency Group::
This parameter specifies the consistency group from which the volumes are
detached.

``-parent-consistency-group <text>`` - Parent Consistency Group::
This parameter specifies the parent consistency group from which the
```

```
volumes are detached.

``-volume <text>`` - Volume::
This parameter specifies the volumes which are reassigned.

{ ``-new-consistency-group <text>`` - New Consistency Group::
This optional parameter specifies a new child consistency group to which
the volumes are reassigned and attached to.

| ``-destination-consistency-group <text>`` - Destination Consistency
Group }::
This optional parameter specifies existing child consistency group to
which the volumes are reassigned and attached to.

== Examples

The following example will reassign volumes vol1,vol2 from child
consistency group named cg in parent consistency group parentCg on a
Vserver named vs0 to new child consistency group new_cg.
```

cluster1::> vserver consistency-group volume reassign -vserver vs0 -consistency-group cg -parent-consistency -group parentCg -volume vol1,vol2 -new-consistency-group new_cg
(vserver consistency-group volume reassign)
[Job 100] Job succeeded: Success

```
[[ID023bdf48c7831fb56587238d370669a3]]
= vserver consistency-group volume remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a volume from consistency group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver consistency-group volume remove` removes volumes from a
consistency group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the Vserver that contains the consistency group
from which the volume will be removed.

``-consistency-group <text>`` - Consistency Group::
This parameter specifies the consistency group from which the volume will
be removed.

``[-parent-consistency-group <text>]`` - Parent Consistency Group::
This parameter specifies the parent consistency group from which the
volume will be removed.

``-volume <volume name>`` - Volume::
This parameter specifies the volume to be removed from the consistency
group.

== Examples

The following example will remove volumes vol1,vol2 from child consistency
group named cg in parent consistency group parentCg on a Vserver named
vs0.
```

cluster1::> vserver consistency-group volume remove -vserver vs0 -consistency-group cg -parent-consistency -group parentCg -volume vol1
(vserver consistency-group volume remove)
Warning: Are you sure you want to remove volume "vol1" from consistency group "cg" in Vserver "vs0" ? {y|n}:

y
[Job 100] Job succeeded: Success

```
[[IDfb7da3271c351e32711f285862c034b3]]
= vserver consistency-group volume show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of existing consistency group Volumes

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The command displays information for consistency group volumes. Use the `
instance ` parameter to display additional consistency group volume
details.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver Name::
Selects information about the consistency group volumes in the specified
Vserver.

``[-consistency-group <text>]`` - Consistency Group::
Selects information about the consistency group volumes in the specified
consistency group.

``[-parent-consistency-group <text>]`` - Parent Consistency Group::
Selects information about the consistency group volumes in the specified
parent consistency group.

``[-volume <volume name>]`` - Volume::
Selects information about the specified consistency group volume.
```

```
``[-uuid <UUID>]`` - Volume UUID::
Selects information about the consistency group volume that matches the
specified volume UUID.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
Selects information about the consistency group volumes that matches the
specified Vserver UUID.

``[-consistency-group-uuid <UUID>]`` - UUID::
Selects information about the consistency group volumes that matches the
specified consistency group UUID.

``[-parent-consistency-group-uuid <UUID>]`` - Parent CG UUID::
Selects information about the consistency group volumes that matches the
specified parent consistency group UUID.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size::
Selects information about the consistency group volumes that have the
specified size.

``[-space-available {<integer>[KB|MB|GB|TB|PB]}]`` - Available Space::
Selects information about the consistency group volumes that have the
specified available space.

``[-space-used {<integer>[KB|MB|GB|TB|PB]}]`` - Used Space::
Selects information about the consistency group volumes that have the
specified used space.

``[-export-policy <export policy name>]`` - Volume Export Policy::
Selects information about the consistency group volumes that have the
specified export policy.

``[-cifs-shares <Share>,...]`` - Volume CIFS Share Names::
Selects information about the consistency group volumes that have the
specified CIFS share name.

== Examples

The following command lists all the volumes that are associated with a
consistency group.
```

```
cluster1::> vserver consistency-group volume show
(vserver consistency-group volume show)
                    Parent
         Consistency  Consistency                    Available  Used
Vserver   Group        Group        Volume       Size   Space     Space
--------- ------------ ------------ ------------ ------- --------- -------
```

```
san_vs0   Child_CG_1   Parent_cg   vol_child1_1  206MB   205.7MB   296KB
san_vs0   Child_CG_2   Parent_cg   vol_child2_1  206MB   205.7MB   280KB
san_vs0   Child_CG_3   Parent_cg   vol_child3_1  206MB   205.7MB   260KB
san_vs0   Child_CG_3   Parent_cg   vol_child3_2  206MB   205.7MB   332KB
san_vs0   Child_CG_3   Parent_cg   vol_child3_3  206MB   205.7MB   296KB
5 entries were displayed.
```

```
:leveloffset: -1


= vserver export-policy commands

:leveloffset: +1


[[ID8a51aee955f8b206b9f763faaf9cb6f9]]
= vserver export-policy check-access
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Given a Volume And/or a Qtree, Check to See If the Client Is Allowed
Access

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy check-access` command checks whether a specific
client is allowed access to a specific export path. This enables you to
test export policies to ensure they work as intended and to troubleshoot
client access issues.

The command takes the volume name (and optionally the qtree name) as input
and computes the export path for the volume/qtree. It evaluates the export
policy rules that apply for each path component and displays the policy
name, policy owner, policy rule index and access rights for that path
component. If no export policy rule matches the specified client IP
address access is denied and the policy rule index will be set to 0. The
output gives a clear view on how the export policy rules are evaluated and
helps narrow down the policy and (where applicable) the specific rule in
the policy that grants or denies access.

== 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.

``-vserver <vserver name>`` - Vserver Name::
This parameter specifies the name of the Vserver in which the export
policy resides.

``-volume <volume name>`` - Volume Name::
This parameter specifies the name of the volume that you want to check
export access for. To check export access for a qtree use the -qtree
parameter. The -qtree parameter is optional. If you specify the -qtree
parameter, you must provide the name of the volume containing the qtree.
If you do not specify the -qtree parameter, export access will be checked
only for the volume.

``-client-ip <IP Address>`` - Client IP Address::
This parameter specifies the IP address of the client that you want to
check export access for.

``-authentication-method <authentication method>`` - Authentication
Method::
This parameter specifies the authentication method of the client that is
attempting access. Possible values include the following:
+
  * ``_sys_`` - The authentication method used by the client is AUTH_SYS.
  * ``_krb5_`` - The authentication method used by the client is Kerberos
v5.
  * ``_krb5i_`` - The authentication method used by the client is Kerberos
v5 with integrity service.
  * ``_krb5p_`` - The authentication method used by the client is Kerberos
v5 with privacy service.
  * ``_ntlm_`` - The authentication method used by the client is CIFS
NTLM.
  * ``_none_`` - The authentication method used by the client is not
explicitly listed in the list of values in the rorule.



``-protocol <Client Access Protocol>`` - Protocol::
```

```
This parameter specifies the protocol that the client is using when
attempting to access the exported path. Possible values include the
following:
+
  * ``_nfs3_`` - The NFSv3 protocol
  * ``_nfs4_`` - The NFSv4 protocol
  * ``_cifs_`` - The CIFS protocol



``-access-type {read|read-write|denied}`` - Access Rights to Check for::
This parameter specifies the type of access you want to check for.
Possible values are read for read-only access and read-write for read-
write access.

``[-qtree <qtree name>]`` - Name of the Qtree::
This optional parameter specifies the qtree in the volume that is part of
the exported path. If you specify this parameter, you must also provide
the name of the volume the qtree belongs to.

``[-path <text>]`` - Path::
Selects the entries in the output that match the specified path value.
This field describes the junction-path path component encountered when
evaluating the export policies starting from the root ('/') of the
Vserver.

``[-policy <text>]`` - Export Policy::
Selects the entries in the output that match the specified policy value.
This field describes the export policy that is in effect for the path
encountered so far when evaluating the export policies starting from the
root ('/') of the Vserver.

``[-policy-owner <text>]`` - Export Policy Owner::
Selects the entries in the output that match the specified policy owner
value. This field describes the owner of the export policy that is in
effect for the path encountered so far when evaluating the export policies
starting from the root ('/') of the vserver. The owner of the export
policy could be a volume or a qtree.

``[-policy-owner-type {volume|qtree}]`` - Type of Export Policy Owner::
Selects the entries in the output that match the specified type of the
owner of an export policy. Possible values include the following:
+
  * ``_volume_`` - The owner of the export policy is a volume
  * ``_qtree_`` - The owner of the export policy is a qtree
```

```
``[-rule-index <integer>]`` - Export Policy Rule Index::
Selects the entries in the output that match the specified export policy
rule index. This field describes the rule index of the rule in the export
policy that grants or denies access. If the value of the rule index is 0
it implies none of the client match strings provided in the rules of the
export policy matched the specified IP address of the client.

``[-access {read|read-write|denied}]`` - Access Rights::
Selects the entries in the output that match the specified access value.
This field describes the access rights to the path. Possible values
include the following:
+
  * ``_read_`` - Read access is granted
  * ``_read-write_`` - Read-write access is granted
  * ``_denied_`` - Requested access is denied



``[-partial-rule-match {true|false}]`` - Did a Subset of the Rules
Match?::
Selects the entries in the output that match if a partially matched subset
of rules in the export policy were used to grant access to the client.

``[-clientmatch <text>]`` - Client Match Spec::
Selects the entries in the output that match the specified clientmatch
string. The clientmatch string denotes the string that resulted in a rule
match for the specified client IP address.

``[-security-style <security style>]`` - Security Style::
Selects the entries in the output that match the specified security style
value. Possible values are unix, ntfs and mixed.

== Examples

The following examples of the `vserver export-policy check-access` command
display various possible results for client export access checks.
```

cluster1::> vserver export-policy check-access -vserver vs1 -client-ip 10.22.32.42 -volume flex_vol
-authentication-method sys -protocol nfs3 -access-type read
Policy Policy Rule Security
Path Policy Owner Owner Type Index Access Style
--- --- --- --- --- --- ---
/ default vs1_root volume 1 read mixed
/dir1 default vs1_root volume 1 read mixed
/dir1/dir2 default vs1_root volume 1 read mixed
/dir1/dir2/flex1 data flex_vol volume 10 read mixed

```
4 entries were displayed.

cluster1::> vserver export-policy check-access -vserver vs1 -client-ip 10.22.32.42 -volume flex_vol
-authentication-method sys -protocol nfs3 -access-type read-write
Policy Policy Rule Security
Path Policy Owner Owner Type Index Access Style
-------------------- ---------- --------- ---------- ---- ---------- ----------
/ default vs1_root volume 1 read mixed
/dir1 default vs1_root volume 1 read mixed
/dir1/dir2 default vs1_root volume 1 read mixed
/dir1/dir2/flex1 data flex_vol volume 10 read-write mixed
4 entries were displayed.

cluster1::> vserver export-policy check-access -vserver vs1 -client-ip 10.22.32.42 -volume flex_vol
-authentication-method sys -protocol nfs3 -access-type read-write -qtree qt1
Policy Policy Rule Security
Path Policy Owner Owner Type Index Access Style
-------------------- ---------- --------- ---------- ---- ---------- ----------
/ default vs1_root volume 1 read mixed
/dir1 default vs1_root volume 1 read mixed
/dir1/dir2 default vs1_root volume 1 read mixed
/dir1/dir2/flex1 data flex_vol volume 10 read mixed
/dir1/dir2/flex1/qt1 primarynames
qt1 qtree 0 denied mixed
5 entries were displayed.

cluster1::> vserver export-policy check-access -vserver vs1 -client-ip 10.22.32.42 -volume flex_vol
-authentication-method ntlm -protocol cifs -access-type read-write -qtree qt1
Policy Policy Rule Security
Path Policy Owner Owner Type Index Access Style
-------------------- ---------- --------- ---------- ---- ---------- ----------
/ default vs1_root volume 1 read mixed
/dir1 default vs1_root volume 1 read mixed
/dir1/dir2 default vs1_root volume 1 read mixed
/dir1/dir2/flex1 data flex_vol volume 10 read mixed
/dir1/dir2/flex1/qt1 primarynames
qt1 qtree 2 denied mixed
5 entries were displayed.
```

```
[[ID0be6d5b8da6be671ba477d85a55a0eaa]]
= vserver export-policy copy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Copy an export policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy copy` command creates a copy of an export
policy on the same or a different Vserver. The command fails if an export
policy with the specified new name already exists on the target Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy that you
want to copy is located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the export policy that you want to copy.

``-newvserver <vserver name>`` - New Vserver::
This parameter specifies the Vserver to which you want to copy the export
policy.

``-newpolicyname <export policy name>`` - New Export Policy Name::
This parameter specifies the name of the new policy.

== Examples

The following example copies an existing policy named read_only_expolicy
located on a Vserver named vs0 to a new policy named default_expolicy
located on a Vserver named vs1.
```

vs1::> vserver export-policy copy -vserver vs0 -policyname read_only_expolicy -newvserver vs1 -newpolicyname default_expolicy

```
[[ID6d1ff8a79946a17ed0d847b4c1071ebf]]
= vserver export-policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a rule set

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy create` command creates an export policy. You
can use the xref:{relative_path}vserver-export-policy-rule-
create.html[vserver export-policy rule create] command to add rules to a
policy. Each cluster has an empty default export policy with the ID 0.
This default export policy does not contain any rules. You cannot delete
the default export policy, but you can rename or modify it.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to create the
export policy.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the export policy that you want to create.

== Examples

The following example creates an export policy named read_only_expolicy on
a Vserver named vs0:
```

vs1::> vserver export-policy create -vserver vs0 -policyname read_only_expolicy

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-create.html[vserver
export-policy rule create]




[[ID9c79cae04ee69b3e320533f5fe23300b]]
= vserver export-policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a rule set

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy delete` command deletes an export policy. You
cannot delete the default policy (named `default` ) for a Vserver unless
you delete the Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy that you
want to delete is located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the export policy that you want to delete.

== Examples

The following example deletes an export policy named test_expolicy from a
Vserver named vs0:
```

vs1::> vserver export-policy delete -vserver vs0 -policyname test_expolicy

```
[[IDa0ae786ca6d4250c36e46713d8e2df23]]
= vserver export-policy rename
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rename an export policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rename` command renames an export policy.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy is
located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the export policy that you want to rename.

``-newpolicyname <export policy name>`` - New Export Policy Name::
This parameter specifies the new name of the export policy.

== Examples

The following example renames an export policy named user_expolicy with
the name read_only_expolicy on a Vserver named vs0:
```

vs1::> vserver export-policy rename -vserver vs0 -policyname user_expolicy -newpolicyname read_only_expolicy

```
[[ID62ec09bbded40c43ba147b2fd0499cad]]
= vserver export-policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of rule sets

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy show` command displays the following
information:

* Vserver name
* Export policy name
* Policy ID (diagnostic privilege level only)



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays a list of export
policies that are located on the Vserver that you specify.

``[-policyname <export policy name>]`` - Policy Name::
If you specify this parameter, the command displays only the export policy
or sets that match the specified name.

== Examples

The following example displays a list of all export policies:
```

vs1::> vserver export-policy show
VServer Policy Name
--- ---
vs0 default_expolicy
vs0 read_only_expolicy
vs1 default_expolicy
vs1 test_expolicy
4 entries were displayed.

```
[[ID53228c12756598fc6e329ca0553460c9]]
= vserver export-policy access-cache flush
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Flush an entry from the access cache

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy access-cache flush` command can be used to
remove all entries in the access cache that belong to the specified export
policy. The command can also be used to remove the access cache entry for
a specific IP address belonging to an export policy. You must provide the
name of the node that hosts the access cache and the name of the Vserver
that owns the export policy. This command differs from the
xref:{relative_path}vserver-export-policy-cache-flush.html[vserver export-
policy cache flush] command. The xref:{relative_path}vserver-export-
policy-cache-flush.html[vserver export-policy cache flush] command allows
you to flush all access cache entries across all export policies in a
Vserver. In contrast the `vserver export-policy access-cache flush`
command gives you the granularity to flush a specific access cache entry
or the granularity to flush all access cache entries for a specific export
policy.

This command is useful to clear out a negative access cache entry. A
negative cache entry is one where a client IP address experiences an
access denied error due to stale export policy rule information present in
the cache entry. ONTAP maintains several caches in the kernel and
userspace to speed access to exports. A negative cache entry can get
created in the access cache if a client tries to access an export path
before the export rules or the name server settings or the caches in
management gateway have been updated to grant access to that client. The
```

```
negative cache entry will remain in the access cache until the TTL for the
entry expires and the entry is refreshed. You can use the `export-policy
access-cache config show ` command to find out the refresh intervals and
timeouts for the access cache. If you know that the caches in userspace
have the latest information for the client and don't want to wait until
the TTL for the access cache entry expires then you can use this command
to remove the access cache entry in the kernel and force the cache entry
to get re-populated with the latest information that will allow the client
to access the export path.



You can use the `vserver export-policy access-cache entry show` and
`vserver export-policy access-cache entry show-rules` commands to examine
the contents of an entry in the access cache before removing it using the
flush command.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
flush the access cache entry.

``-node <nodename>`` - Node::
This parameter specifies the node on which you want to flush the access
cache entry.

``-policy <text>`` - Export Policy Name::
This parameter specifies the name of the export policy that is effective
for the exported path that the client is trying to access.

``[-address <IP Address>]`` - IP Address::
This parameter is optional. It specifies the IP address of the client
whose access cache entry you want to remove. If this parameter is not
specified all access cache entries belonging to the specified export
policy will be removed.

== Examples

The following example flushes the access cache entry for client IP address
'1.2.3.4' in volume 'flex1' having export policy 'testpol' in a Vserver
named 'vs1' on node 'vsim1':
```

cluster1::*> vserver export-policy access-cache flush -vserver vs1 -node vsim1 -policy testpol -address 1.2.3.4
Successfully removed access cache entry for IP address "1.2.3.4" belonging to export policy "testpol" in Vserver "vs1" on node "vsim1".

cluster1::*> vserver export-policy access-cache flush -vserver vs1 -node vsim1 -policy testpol

Warning: This command removes all access cache entries for export policy "testpol" in Vserver "vs1" on node "vsim1". Do you want to continue? {y|n}: y

Successfully removed 1 access cache entry for export policy "testpol" in Vserver "vs1" on node "vsim1".

```
== Related Links

* xref:{relative_path}vserver-export-policy-cache-flush.html[vserver
export-policy cache flush]




[[IDd80d4110a541a22b64474420460fdeeb]]
= vserver export-policy access-cache show-negative
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about the negative access cache entry

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver export-policy access-cache show-negative` command can be used
to display the contents of access cache having negative polarity of the
specified node for a particular client IP address belonging to an export
policy in a Vserver.

The command will display information such as the age, policy name and
client IP address of the negative access cache entry.



If you are interested in finding out more details about the access cache
then you can use the xref:{relative_path}vserver-export-policy-access-
cache-show.html[vserver export-policy access-cache show] command.

If the client IP address for which access is denied is not cached in the
access cache then the command will display an error message stating that
this table is current empty.
```

```
== 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>|local}]`` - Node Name::
This parameter specifies the node on which you want to examine the access
cache entry.Use an exact value. Queries are not supported.

``[-vserver <vserver>]`` - Vserver::
This parameter specifies the name of the Vserver on which you want to see
the access cache entry. Use an exact value. Queries are not supported.

``[-policy <text>]`` - Export Policy Name::
This parameter specifies the name of the export policy that is in effect
on the export path that the client is trying to access.Use an exact value.
Queries are not supported.

``[-client-ip <IP Address>]`` - Client IP Address::
This parameter specifies the IP address of the client whose access cache
entry you want to examine.Use an exact value. Queries are not supported.

``[-age <[<integer>h][<integer>m][<integer>s]>]`` - Age of Entry::
Selects the access cache entries that match the specified age of the
entry. This field describes the age of the access cache entry.

== Examples

The following example shows the contents of the access cache entry having
negative polarity:
```

cluster1::*> vserver export-policy access-cache show-negative
Node: vikash2-vsim1
Vserver: vs12
Policy Name IP Address Age of Entry
---------- --------- ----------
default 1.1.1.1 16s
default 1.1.1.2 17s
default 1.1.1.3 18s
3 entries were displayed.

```
== Related Links

* xref:{relative_path}vserver-export-policy-access-cache-show.html[vserver
export-policy access-cache show]




[[IDd9b26416dbb0316bfa5b32fb12b50a44]]
= vserver export-policy access-cache show-rules
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about the export policy rules in the access cache
entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy access-cache show-rules` command is used in
conjunction with the xref:{relative_path}vserver-export-policy-access-
cache-show.html[vserver export-policy access-cache show] command. The
xref:{relative_path}vserver-export-policy-access-cache-show.html[vserver
export-policy access-cache show] command displays the state and contents
of an access cache entry on the specified node for a particular client IP
address belonging to an export policy in a Vserver. The command lists the
rule indexes of the export policy rules that matched. If you are
interested in finding out the security settings for each policy rule that
matched then you can use the ` vserver export-policy access-cache show-
rules` command. You can use the -instance switch to get a more detailed
listing. Do note that the security settings of the rules cached in the
access cache entry match the security settings of the rules that can be
obtained by running the xref:{relative_path}vserver-export-policy-rule-
show.html[vserver export-policy rule show] command with the corresponding
rule index.

If the client IP address is not cached in access cache then the command
will display an error message stating that the entry does not exist.

== 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 specifies the node on which you want to examine the export
policy rule details in the access cache entry.

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to see
the policy rule details in the access cache entry.

``-policy <export policy name>`` - Policy Name::
This parameter specifies the name of the export policy that is in effect
on the export path that the client is trying to access.

``-address <IP Address>`` - IP Address::
This parameter specifies the IP address of the client whose access cache
entry you want to examine in greater detail.

``[-ruleindex <integer>]`` - Entry Policy Rule Index::
This optional parameter specifies the index number of the export rule of a
specific policy.

``[-protocol <Client Access Protocol>,...]`` - Access Protocol::
This optional parameter specifies the list access protocols of export
rules.

``[-rorule <authentication method>,...]`` - RO Access Rule::
This parameter specifies the security type for read-only access to volumes
that use the export rule.

``[-rwrule <authentication method>,...]`` - RW Access Rule::
This parameter specifies the security type for read-write access to
volumes that use the export rule.

``[-superuser <authentication method>,...]`` - Superuser Security Types::
This parameter specifies a security type for superuser access to files.

``[-anon-uid <integer>]`` - Anonymous User ID::
This parameter specifies an anonymous user ID that the user credentials
are mapped to.
```

```
``[-anon-gid <integer>]`` - Anonymous User Primary GID::
This parameter specifies an anonymous User Primary GID.

``[-anon-gid-list <integer>,...]`` - Anonymous User GID List::
This parameter specifies an anonymous User Primary GID list.

``[-protocol-flags {allow-suid|allow-dev|allow-nfs-tls-only}]`` - Protocol
Flags::
This parameter specifies protocol flags such as allow-suid, allow-dev and
allow-nfs-tls-only.

``[-ntfs-unix-security-ops {ignore|fail}]`` - NTFS Unix Security Options::
This parameter specifies whether UNIX-type permissions changes on NTFS
(Windows) volumes are prohibited (fail) or allowed (ignore).

``[-chown-mode {restricted|unrestricted}]`` - Change Ownership Mode::
This parameter specifies a change ownership mode.

``[-clientmatch <text>]`` - Client Match String::
This parameter specifies the client or clients to which the export rule
applies.

``[-anonuser <text>]`` - Anonymous Username or ID::
This parameter specifies a UNIX user ID or user name that the user
credentials are mapped to.

== Examples

The following example shows the contents of the access cache entry for
client IP address '1.2.3.4' in volume 'flex1' having export policy
'testpol' in a Vserver named 'vs1' on node 'vsim1'. This entry has two
export policy rules with rule indexes 1 and 2 that matched and are cached
in the entry. To examine what the rule settings are in each of these rules
we can use the show-rules variant of the command.
```

cluster1::*>vserver export-policy access-cache show -vserver vs1 -node vsim1 -policy testpol -address 1.2.3.4
Node: vsim1
Vserver: vs1
Policy Name: testpol
IP Address: 1.2.3.4
Access Cache Entry Flags: -
Result Code: 0
Failure Type Code: 0
Number of Matched Policy Rules: 2
List of Matched Policy Rule Indexes: 1, 2
Age of Entry: 5s

```
cluster1::*>vserver export-policy access-cache show-rules -vserver vs1 -node vsim1 -policy testpol -address
1.2.3.4
Rule Access RO RW Super Anon
Node Address Policy Index Protocol Rule Rule User User
---------- --------------- --------- ----- -------- ---- ---- ----- -----
vsim1 1.2.3.4 testpol 1 any any any none 65534
vsim1 1.2.3.4 testpol 2 nfs3 never never sys 123
2 entries were displayed.

cluster1::*>vserver export-policy access-cache show-rules -vserver vs1 -node vsim1 -policy testpol -address
1.2.3.4 -instance
Vserver: vs1
Node: vsim1
Policy Name: testpol
IP Address: 1.2.3.4
Export Policy ID: 12884901890
Entry Policy Rule Index: 1
Access Protocol: any
RO Access Rule: any
RW Access Rule: any
Superuser Security Types: none
Anonymous User ID: 65534
Protocol Flags: allow-suid, allow-dev
NTFS Unix Security Options: fail
Change Ownership Mode: restricted
Vserver: vs1
Node: vsim1
Policy Name: testpol
IP Address: 1.2.3.4
Export Policy: testpol
Export Policy ID: 12884901890
Entry Policy Rule Index: 2
Access Protocol: nfs3
RO Access Rule: never
RW Access Rule: never
Superuser Security Types: sys
Anonymous User ID: 123
Protocol Flags: allow-suid
NTFS Unix Security Options: ignore
Change Ownership Mode: restricted
2 entries were displayed.

cluster1::*> vserver export-policy rule show -vserver vs1 -policyname testpol -ruleindex 1
Vserver: vs1
Policy Name: testpol
Rule Index: 1
Access Protocol: any
Client Match Hostname, IP Address, Netgroup, or Domain: 0.0.0.0/0
RO Access Rule: any
RW Access Rule: any
User ID To Which Anonymous Users Are Mapped: 65534
Superuser Security Types: none
Honor SetUID Bits in SETATTR: true
Allow Creation of Devices: true
```

cluster1::*> vserver export-policy rule show -vserver vs1 -policyname testpol -ruleindex 2
Vserver: vs1
Policy Name: testpol
Rule Index: 2
Access Protocol: nfs3
Client Match Hostname, IP Address, Netgroup, or Domain: 0.0.0.0/0
RO Access Rule: never
RW Access Rule: never
User ID To Which Anonymous Users Are Mapped: testu1
Superuser Security Types: sys
Honor SetUID Bits in SETATTR: true
Allow Creation of Devices: false

```
== Related Links

* xref:{relative_path}vserver-export-policy-access-cache-show.html[vserver
export-policy access-cache show]
* xref:{relative_path}vserver-export-policy-rule-show.html[vserver export-
policy rule show]




[[ID374de2c5cdadb342ac8fddc97892e8aa]]
= vserver export-policy access-cache show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about the access cache entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy access-cache show` command can be used to
display the contents of an access cache entry of the specified node for a
particular client IP address belonging to an export policy in a Vserver.

The command will display information such as the flags of the access cache
entry, the age of the entry, any errors that were encountered when looking
up the export policy rules from the management gateway, and the number of
policy rules from the export policy that matched the specified client IP
address. If an error is encountered when looking up the export policy
rules from the management gateway process, the first rule index in the
```

```
export policy that encountered the error is displayed. The client match
string or the anon string in the rule that caused the rule evaluation to
fail is also displayed. A more detailed view of the output of this command
is available if you specify the -instance switch to the command.



The command output lists the rule indexes of the policy rules that
matched. If you are interested in finding out the security settings for
each policy rule that matched then you can use the
xref:{relative_path}vserver-export-policy-access-cache-show-
rules.html[vserver export-policy access-cache show-rules] command.

If the client IP address is not cached in the access cache then the
command will display an error message stating that the entry does not
exist.

== 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 specifies the node on which you want to examine the access
cache entry.

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to see
the access cache entry.

``-policy <export policy name>`` - Policy Name::
This parameter specifies the name of the export policy that is in effect
on the export path that the client is trying to access.

``-address <IP Address>`` - IP Address::
This parameter specifies the IP address of the client whose access cache
entry you want to examine.

``[-flags {pending|refreshing|is-abandoned|is-queued-for-update|is-
updating|has-usable-data}]`` - Access Cache Entry Flags::
Selects the access cache entries that match the specified flags value. The
```

```
flags describe the internal state of the access cache entry. The access
cache entry could be in 'pending' state. This denotes the initial state of
the access cache entry when a client first tries to access the exported
mount point and the rules in the export policy are being matched against
the IP address of the client. The 'refreshing' state denotes that the
access cache entry is being refreshed. The 'abandoned' state denotes that
the access cache entry has been cleared as a result of a cache flush
operation. If the access cache entry has been successfully evaluated this
field will not be set to any value.

``[-result <integer>]`` - Result Code::
Selects the access cache entries that match the specified result value.
This field describes the error code of the error encountered when matching
the IP address of the client against the rules specified in the export
policy. If all rules were evaluated successfully this field will be set to
0.

``[-first-unresolved-index <integer>]`` - First Unresolved Rule Index::
Selects the access cache entries that match the specified unresolved rule
index value. This field describes the rule index of the first rule in the
export policy that could not be evaluated successfully when matching the
IP address of the client against the rules specified in the export policy.
If all rules were evaluated successfully this field will not be set to any
value.

``[-unresolved-clientmatch <text>]`` - Unresolved Clientmatch::
Selects the access cache entries that match the specified unresolved
client match value. This field describes the client match string that
caused the rule evaluation to fail at the displayed rule index. Client
match strings that denote a netgroup, hostname or a domain name can fail
in evaluation if there are problems in contacting the name servers
configured to serve them. If all rules were evaluated successfully this
field will not be set to any value.

``[-num-rules <integer>]`` - Number of Matched Policy Rules::
Selects the access cache entries that match the specified number of
matched export rules. This field describes the number of rules in the
export policy that were matched successfully against the IP address of the
client. If the number of matched rules is 0 and the 'result' field is also
0 then the client will experience an access denied error during mount. If
the number of matched rules is non-zero and the 'result' field is 0 then
access is granted or denied based on the ro, rw, superuser and other
security settings in the matched rules. If the number of matched rules is
0 and the 'result' field has a non-zero value in it the client will
experience a hang until the error that caused the rule evaluation to fail
is resolved. If the number of matched rules is non-zero and the 'result'
```

```
field has a non-zero value then this represents a situation where an error
was encountered that stopped the match of rules in the export policy
against the IP address of the client. The rules that have matched so far
are used to make access decisions. (Note that the match of rules follows
an ordering precedence determined by the rule index). Access may be
granted if the security settings in the rules that have matched so far
allow access. The security settings in the partial subset of matched rules
are never used to deny access because they represent an incomplete set of
matched export rules. Instead the client will experience a hang until the
error that caused the rule evaluation to fail is resolved.

``[-ruleindex-list <integer>,...]`` - List of Matched Policy Rule
Indexes::
Selects the access cache entries that match the specified list of matched
rule indexes. This field describes a comma separated list of the indexes
of the rules in the export policy that matched the IP address of the
client. If no rules match the IP address of the client or an error was
encountered in the client match process then this field will not be set to
any value.

``[-age <[<integer>h][<integer>m][<integer>s]>]`` - Age of Entry::
Selects the access cache entries that match the specified age of the
entry. This field describes the age of the access cache entry.

``[-polarity {positive|negative|init}]`` - Access Cache Entry Polarity::
Selects the access cache entries that match the specified polarity of the
entry. The polarity of an access cache entry can be positive or negative.
A positive polarity denotes that access is granted to the client IP
address. A negative polarity denotes that access is denied to the client
IP address.

``[-duration-since-last-use <[<integer>h][<integer>m][<integer>s]>]`` -
Time Elapsed since Last Use for Access Check::
Selects the access cache entries that match the specified time duration
since the entry was last used for access determination.

``[-duration-since-last-update-attempt
<[<integer>h][<integer>m][<integer>s]>]`` - Time Elapsed since Last Update
Attempt::
Selects the access cache entries that match the specified time duration
since the access cache entry was last updated.

``[-last-update-attempt-result <integer>]`` - Result of Last Update
Attempt::
Selects the access cache entries that match the specified result obtained
when the access cache entry was last updated.
```

```
``[-clientmatch-list <text>,...]`` - List of Client Match Strings::
Selects the access cache entries that match the specified list of
clientmatch strings that matched the specified client IP address.

== Examples

The following example shows the contents of the access cache entry for
client IP address '10.22.33.32' in volume 'flex1' having export policy
'testpol' in a Vserver named 'vs1' on node 'vsim1':
```

cluster1::*> vserver export-policy access-cache show -vserver vs1 -policy testpol -node vsim1 -address 10.22.33.32
Node: vsim1
Vserver: vs1
Policy Name: testpol
IP Address: 10.22.33.32
Access Cache Entry Flags: has-usable-data
Result Code: 0
First Unresolved Rule Index: -
Unresolved Clientmatch: -
Number of Matched Policy Rules: 1
List of Matched Policy Rule Indexes: 20
Age of Entry: 77s
Access Cache Entry Polarity: positive
Time Elapsed since Last Update Attempt: 8s
Time Elapsed since Last Use for Access Check: 3s
Result of Last Update Attempt: 7208
List of Client Match Strings: 0.0.0.0/0

```
== Related Links

* xref:{relative_path}vserver-export-policy-access-cache-show-
rules.html[vserver export-policy access-cache show-rules]




[[IDc21ed77f8662502325d937b48e38ccf2]]
= vserver export-policy access-cache config modify-all-vservers
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify exports access cache configuration for all Vservers

*Availability:* This command is available to _cluster_ administrators at
```

```
the _admin_ privilege level.

== Description

The `vserver export-policy access-cache config modify-all-vservers`
command modifies access cache timeout values for all Vservers. Modifying
these values from any node updates the values on all the nodes in the
cluster. The modified values persist across reboots.

NOTE: This command is not supported in a cluster with effective cluster
version of ONTAP 9.0.0 or later. The access cache settings are modified on
a per-Vserver basis starting ONTAP 9.0.0. See the
xref:{relative_path}vserver-export-policy-access-cache-config-
modify.html[vserver export-policy access-cache config modify] command.



== Parameters

``[-ttl-positive <integer>]`` - TTL For Positive Entries (Secs)::
This parameter specifies the duration after which positive access cache
entries will be refreshed when the client accesses.

``[-ttl-negative <integer>]`` - TTL For Negative Entries (Secs)::
This parameter specifies the duration after which negative access cache
entries will be refreshed when the client accesses.

``[-harvest-timeout <integer>]`` - Harvest Timeout (Secs)::
This parameter specifies the time period after which ONTAP deletes unused
entries in the access cache.

``[-isDnsTTLEnabled {true|false}]`` - Is Dns TTL Enabled::
This parameter specifies the dns TTL is enabled or not.

== Examples

The following command sets the positive TTL value to 36000 seconds, the
negative TTL value to 3600 seconds, and the harvest timeout value to 43200
seconds for all Vservers in a cluster where the effective cluster version
is earlier than ONTAP 9.0.0.
```

cluster1::*> vserver export-policy access-cache config modify-all-vservers -ttl-positive 36000 -ttl-negative 3600 -harvest-timeout 43200 -isDnsTTLEnabled false

cluster1::*> vserver export-policy access-cache config show-all-vservers
TTL For Positive Entries (secs): 36000
TTL For Negative Entries (secs): 3600

Harvest Timeout (secs): 43200
Is Dns TTL Enabled: false

```
== Related Links

* xref:{relative_path}vserver-export-policy-access-cache-config-
modify.html[vserver export-policy access-cache config modify]




[[ID29ba3ef22d399e3dbf7d81a0828adf73]]
= vserver export-policy access-cache config modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify exports access cache configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy access-cache config modify` command modifies
access cache timeout values per Vserver. Modifying these values from any
node updates the values on all the nodes in the cluster. The modified
values persist across reboots.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver name for which the timeout values
need to be modified.

``[-ttl-positive <integer>]`` - TTL For Positive Entries (Secs)::
This parameter specifies the duration after which positive access cache
entries will be refreshed upon client access. The value is specified in
seconds. The default value is 3600 seconds. Valid values range from 300
seconds to 86400 seconds.

``[-ttl-negative <integer>]`` - TTL For Negative Entries (Secs)::
This parameter specifies the duration after which negative access cache
entries will be refreshed upon client access. The value is specified in
```

```
seconds. The default value is 3600 seconds. Valid values range from 60
seconds to 86400 seconds.

``[-harvest-timeout <integer>]`` - Harvest Timeout (Secs)::
This parameter specifies the time period after which ONTAP deletes unused
entries in the access cache. The value is specified in seconds. The
default value is 86400 seconds. Valid values range from 60 seconds to
2592000 seconds.

``[-isDnsTTLEnabled {true|false}]`` - Is Dns TTL Enabled::
This parameter specifies the dns TTL is enabled or not. If dns TTL is
enable then access cache will use ttl returned from dns lookup, in case
dns lookup doesn't return TTL then it will use default ttl value.

== Examples

The following command sets the positive TTL value to 36000 seconds, the
negative TTL value to 3600 seconds, and the harvest timeout value to 43200
seconds for Vserver 'vs0':
```

cluster1::*> vserver export-policy access-cache config modify -ttl-positive 36000 -ttl-negative 3600 -harvest -timeout 43200 -isDnsTTLEnabled false

cluster1::*> vserver export-policy access-cache config show -vserver vs0
Vserver: vs0
TTL For Positive Entries (secs): 36000
TTL For Negative Entries (secs): 3600
TTL For Entries with Failure (secs): 1
Harvest Timeout (secs): 43200
Is Dns TTL Enabled: false

```
[[ID02289f2f4ff7daa5212d87f67ea2e4ce]]
= vserver export-policy access-cache config show-all-vservers
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display exports access cache configuration for all Vservers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver export-policy access-cache config show-all-vservers` command
```

```
displays the timeout attributes related to the exports access cache. The
access cache maintains export rules applicable to a client that is
accessing the volume or qtree. ONTAP obtains the access cache timeout
values from the node where you run the command. The command output
displays the following timeout parameters and their values:

* TTL for Positive Entries: This is the TTL for positive entries in the
access cache. During client access, if the TTL for the access cache entry
that is allowing access has expired, that access cache entry will be
refreshed. While the refresh is in progress, client access will be
evaluated with the existing information in the access cache entry.
* TTL for Negative Entries: This is the TTL for negative entries in the
access cache. During client access, if the TTL for the access cache entry
that is denying access has expired, that access cache entry will be
refreshed. While the refresh is in progress, client access will be
evaluated with the existing information in the access cache entry.
* Harvest Timeout: If ONTAP does not use an entry that is stored in the
access cache for this period of time, it deletes the entry.



NOTE: This command is not supported in a cluster with effective cluster
version of ONTAP 9.0.0 or later. The access cache settings are stored on a
per-Vserver basis starting ONTAP 9.0.0. See the
xref:{relative_path}vserver-export-policy-access-cache-config-
show.html[vserver export-policy access-cache config show] command.



== Examples

The following command displays the exports access cache timeout values for
all Vservers in a cluster where the effective cluster version is earlier
than ONTAP 9.0.0:
```

cluster1::*> vserver export-policy access-cache config show-all-vservers
TTL For Positive Entries (secs): 36000
TTL For Negative Entries (secs): 3600
Harvest Timeout (secs): 43200

```
== Related Links

* xref:{relative_path}vserver-export-policy-access-cache-config-
show.html[vserver export-policy access-cache config show]
```

```
[[IDbbcc8fc26f3677c53a39c48be9d75d79]]
= vserver export-policy access-cache config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display exports access cache configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy access-cache config show` command displays the
timeout attributes related to the exports access cache. The access cache
maintains export rules applicable to a client that is accessing the volume
or qtree. The command output displays the following timeout parameters and
their values for each Vserver:

* TTL for Positive Entries: This is the TTL for positive entries in the
access cache. During client access, if the TTL for the access cache entry
that is allowing access has expired, that access cache entry will be
refreshed. While the refresh is in progress, client access will be
evaluated with the existing information in the access cache entry.
* TTL for Negative Entries: This is the TTL for negative entries in the
access cache. During client access, if the TTL for the access cache entry
that is denying access has expired, that access cache entry will be
refreshed. While the refresh is in progress, client access will be
evaluated with the existing information in the access cache entry.
* TTL for Entries with Failure: This is the TTL for access cache entries
for which a failure was encountered while trying to get matching rules.
* Harvest Timeout: If ONTAP does not use an entry that is stored in the
access cache for this period of time, it deletes the entry.



== 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.

``[-vserver <vserver name>]`` - Vserver::
If this parameter is specified, the command displays the timeout values
for the specified Vserver.

``[-ttl-positive <integer>]`` - TTL For Positive Entries (Secs)::
If this parameter is specified, the command displays the timeout values
for Vservers whose ttl-positive matches the provided value.

``[-ttl-negative <integer>]`` - TTL For Negative Entries (Secs)::
If this parameter is specified, the command displays the timeout values
for Vservers whose ttl-negative matches the provided value.

``[-harvest-timeout <integer>]`` - Harvest Timeout (Secs)::
If this parameter is specified, the command displays the timeout values
for Vservers whose harvest-timeout matches the provided value.

``[-isDnsTTLEnabled {true|false}]`` - Is Dns TTL Enabled::
If this parameter is specified, the command displays the isDnsTTLEnabled
value.

== Examples

The following command displays the exports access cache timeout values for
all Vservers in the cluster:
```

cluster1::*> vserver export-policy access-cache config show
Vserver TTL Positive TTL Negative TTL Failure Harvest Timeout isDnsTTLEnabled
(secs) (secs) (secs) (secs)
-------- ------------ ------------ ----------- --------------- --------
vs0 300 60 1 3600 false
vs1 36000 3600 5 3600 false
2 entries were displayed.

```
[[ID66336a5269cf19759f2761f25f5ce66f]]
= vserver export-policy cache flush
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Flush the Export Caches
```

```
*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver export-policy cache flush` command clears out the contents of
the export policy caches for a Vserver. You might need to flush the caches
to allow the changes to immediately take effect for your NFS clients
because of:

* A change to your export policy rules.
* Modifying a host name record in a name server (i.e., local hosts or
DNS).
* Modifying a PTR record in a DNS server (i.e., reverse DNS lookup).
* Modifying the entries in a netgroup in a name server (i.e., local
netgroup, LDAP, or NIS).
* Recovering from a network outage that resulted in a netgroup being
partially expanded.

To flush the caches, you must specify the following items:

* Vserver: either a specific Vserver or use "*" to flush all of them.

You can optionally specify the following items:

* Node: if flushing the ``_access_`` cache, you can also specify which
node to flush it on.
* Cache to flush: by default all but ``_showmount_`` will be flushed.

Note that the ``_showmount_`` cache is not used to determine NFS client
access and as such is only flushable explicitly.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
flush the caches.

``[-node <nodename>]`` - Node::
This parameter specifies the node on which you want to flush the
``_access_`` cache.

``[-cache {all|access|host|id|name|netgroup|showmount|ip}]`` - Cache
Name::
This parameter specifies the name of the cache which you want to flush.
Possible values include the following:
+
```

```
  * ``_all_`` - All caches but ``_showmount_`` . This is the default.
  * ``_access_`` - The export-policy rules access cache.
  * ``_host_`` - The host name to IP cache.
  * ``_id_`` - The ID to credential cache.
  * ``_ip_`` - The IP to host name cache.
  * ``_name_`` - The name to ID cache.
  * ``_netgroup_`` - The netgroup cache.
  * ``_showmount_`` - The showmount caches.



== Examples

The following example flushes the access cache on a Vserver named vs0:
```

cluster1::> vserver export-policy cache flush -vserver vs0 -cache access

```
[[ID7af3835782a9f9fa8b2b00460d0ab4a3]]
= vserver export-policy config-checker show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the status of export policy configuration checker jobs

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver export-policy config-checker show` command displays status
information about export policy configuration checker job. This command
displays the following information:

* Vserver name
* Export policy name
* Export policy configuration checker job state
* Export policy rule checked count
* Export policy rule being checked rule index
* Export policy rule with issue count



NOTE: This command output will only be available after running the export
```

```
policy configuration checker job.



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays export policy
configuration checker job state information for Vservers that match the
specified value.

``[-policy <export policy name>]`` - Policy Name::
If you specify this parameter, the command displays export policy
configuration checker job state information for policy that match the
specified value.

``[-rules-checked <integer>]`` - Number of Rules Checked::
If you specify this parameter, the command displays export policy
configuration checker job state information that have the specified rules-
checked count matching.

``[-rule-being-checked <integer>]`` - Rule Being Checked::
If you specify this parameter, the command displays export policy
configuration checker job state information that have the specified rule-
being-checked index matching.

``[-rules-with-issues <integer>]`` - Number of Rules with Issues::
If you specify this parameter, the command displays export policy
configuration checker job state information that have the specified rules-
with-issues count matching.

``[-state
{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Re
schedule|Error|Quit|Dead|Unknown|Restart|Dormant}]`` - Job State::
If you specify this parameter, the command displays export policy
configuration checker job state information that have the specified state
matching.

== Examples
```

```
The following example displays an export policy configuration checker job
state information for vserver vs2 and policy default:
```

cluster1::> vserver export-policy config-checker show -vserver vs2 -policy default
Job Rules Rule Index Rules With
Vserver Policy State Checked Being Checked Issues
---------- ---------- --------- --------- ---------- ----------
vs2 default Running 1 2 1

```
[[IDeb5dbf35176f8ffa213ffa551513c51c]]
= vserver export-policy config-checker start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start export policy configuration checker job

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy config-checker start` command invokes
background job, which will check export policy configuration and if issue
found in rules then error entry is created for each affected rule in
export policy configuration checker error rule list.

NOTE: Export policy configuration checker only validates hostname,
netgroup and anonymous user related configuration.



== Parameters

``-vserver <vserver name>`` - Vserver::
If you specify this parameter, the export policy configuration checker job
will be triggered for specified Vserver.

``[-policy <export policy name>]`` - Export Policy Name::
If you specify this parameter, the export policy configuration checker job
will be triggered for specified policy.

== Examples

The following example start a export policy configuration checker job for
vserver vs2 and policy default:
```

cluster1::> vserver export-policy config-checker start -vserver vs2 -policy default
[Job 644] Job is queued: Export Policy configuration checker.

```
[[ID3ecc5fef95fa7700cb4ac075325d7c34]]
= vserver export-policy config-checker stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop export policy configuration checker job

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy config-checker stop` command stops running
export policy configuration checker job.

NOTE: Export policy configuration checker stop command only works if the
keys provided are same as the keys provided at the time of starting export
policy configuration checker job.



== Parameters

``-vserver <vserver name>`` - Vserver::
If you specify this parameter, the command stops export policy
configuration checker job, if any export policy configuration checker job
is running for the specified Vserver.

``[-policy <export policy name>]`` - Export Policy Name::
If you specify this parameter, the command stops export policy
configuration checker job, if any export policy configuration checker job
is running for the specified policy.

== Examples

The following example stop an export policy configuration checker job for
Vserver vs2 and policy default:
```

cluster1::> vserver export-policy config-checker stop -vserver vs2 -policy default

```
[[IDe4a46ed8c2c4b8908d76ed49d8a730b4]]
= vserver export-policy config-checker rule delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete error entries for rules from export policy configuration checker
error rule list

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver export-policy config-checker rule delete` command deletes
error rule entries from export policy configuration checker error rule
list. You can delete a specific error entry rule by specifying its rule
index number.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node on which the export policy configuration
error rule entries are stored.

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver which contains the export policy.

``-policy <export policy name>`` - Policy Name::
This parameter specifies the export policy from which you want to delete
an error rule entry.

``-rule-index <integer>`` - Rule Index::
This parameter specifies the index number of the error rule entry that you
want to delete. You can use the xref:{relative_path}vserver-export-policy-
config-checker-rule-show.html[vserver export-policy config-checker rule
show] command to view a list of rules with their index numbers.

== Examples

The following example deletes an error rule entry from config-checker
error rule list, with the index number 1 from an export policy named
default on a Vserver named vs34:
```

cluster1::>vserver export-policy config-checker rule delete -node node-vsim3 -vserver vs34 -policy test -rule -index 1
(vserver export-policy config-checker rule delete)
1 entry was deleted.

```
== Related Links

* xref:{relative_path}vserver-export-policy-config-checker-rule-
show.html[vserver export-policy config-checker rule show]




[[ID2d4663eae2f3ea02e4acd3e90dad25a9]]
= vserver export-policy config-checker rule show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show error entries for rules in export policy configuration checker job

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver export-policy config-checker rule show` command displays
information about error related to configuration in export policy rules.
If a rule has any issues the configuration checker job will log
information about such errors on the node where the job runs. The command
displays the following information:

* Node name
* Vserver name
* Export policy name
* Export policy rule index number
* Export policy rule error



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.
```

```
| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-node {<nodename>|local}]`` - Node::
If you specify this parameter, the command displays detailed error
information for node that matches the specified value.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays detailed error
information for Vservers that match the specified value.

``[-policy <export policy name>]`` - Policy Name::
If you specify this parameter, the command displays detailed error
information for policy that match the specified value.

``[-rule-index <integer>]`` - Rule Index::
If you specify this parameter, the command displays detailed error
information for rule-index that match the specified value.

``[-error <text>]`` - Error Details::
If you specify this parameter, the command displays rule index information
for error that match the specified value. The complete error string needs
to be specified within "{}".

== Examples

The following example displays information about error related to export
rules:
```

cluster1::> vserver export-policy config-checker rule show -node node-vsim3 -vserver vs34 -policy test
(vserver export-policy config-checker rule show)
Rule
Node Vserver Policy Index Error
---------- ------------ --------- --------- ----------
node-vsim3
vs34 test 1 DNS lookup for host "h1" failed
vs34 test 2 Entry not found for "UserName: testuser", DNS lookup for host "h2" failed
2 entries were displayed.

cluster1::> vserver export-policy config-checker rule show -node node-vsim3 -vserver vs34 -policy test -rule -index 1
(vserver export-policy config-checker rule show)
Node: node-vsim3
Vserver: vs34

Policy Name: test
Rule Index: 1
Error Details: DNS lookup for host "h1" failed

```
cluster1::> vserver export-policy config-checker rule show -node node-vsim3 -vserver vs34 -policy test -error
{DNS lookup for host "h1" failed}
(vserver export-policy config-checker rule show)
Rule
Node Vserver Policy Index Error
---------- ------------ --------- --------- ----------------------------------
node-vsim3
vs34 test 1 DNS lookup for host "h1" failed
```

```
[[IDc30327397607c5cf52b278eb173cadb9]]
= vserver export-policy netgroup check-membership
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Check to see if the client is a member of the netgroup

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy netgroup check-membership` command determines
if the client IP address is a member of the netgroup. The netgroup must be
configured as clientmatch in at least one of the export-policy rules
configured in the vserver. ONTAP can determine the membership information
only after it has fully loaded the netgroup into the cache. Until then,
while the reverse lookup scan algorithm might find a match, both DNS round
robin and DNS aliases prevent ruling out non-matches. You can use the
xref:{relative_path}vserver-export-policy-netgroup-queue-show.html[vserver
export-policy netgroup queue show] command to monitor the loading of the
netgroup.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver whose netgroup you want
to check for client membership.

``-netgroup <text>`` - Name of the Netgroup::
This parameter specifies the name of the netgroup that you want to check
for client membership.

``-client-ip <IP Address>`` - Client Address::
This parameter specifies the IP address of the client whose netgroup
membership you want to check.

== Examples

The following examples of the `vserver export-policy netgroup check-
membership` command display various possible results for client membership
checks.
```

cluster1::*> vserver export-policy netgroup check-membership -vserver vs1 -netgroup mercury -client-ip 172.17.16.72
Client 172.17.16.72 is a member of netgroup "mercury" for Vserver "vs1" with state "reverse lookup scan".

cluster1::*> vserver export-policy netgroup check-membership -vserver vs1 -netgroup mercury -client-ip 172.17.16.72
Client 172.17.16.72 is a member of netgroup "mercury" for Vserver "vs1" with state "cache".

cluster1::*> vserver export-policy netgroup check-membership -vserver vs1 -netgroup mercury -client-ip 172.17.16.14
Client 172.17.16.14 is not a member of netgroup "mercury" for Vserver "vs1".

cluster1::*> vserver export-policy netgroup check-membership -vserver vs1 -netgroup big -client-ip 172.17.16.69
Cannot yet determine the membership of client 172.17.16.69 in netgroup "big" for Vserver "vs1". Try again when the netgroup is loaded in the cache.

cluster1::*> vserver export-policy netgroup check-membership -vserver vs1 -netgroup big -client-ip 172.17.16.69
Client 172.17.16.72 is a member of netgroup "big" for Vserver "vs1" with state "cache".

cluster1::*> vserver export-policy netgroup check-membership -vserver vs1 -netgroup big -client-ip 2002:c65f:e228:0:0:0:0:0
Cannot yet determine the membership of client 2002:c65f:e228:: in netgroup "big" for Vserver "vs1". Try again when the netgroup is loaded in the cache.

```
== Related Links

* xref:{relative_path}vserver-export-policy-netgroup-queue-
show.html[vserver export-policy netgroup queue show]




[[IDf4eb6d736d9763e643a63b102e66472d]]
= vserver export-policy netgroup cache show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the Netgroup Cache

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy netgroup cache show` command displays the
```

```
contents of the export policy netgroup cache for a Vserver. Entries shown
here correspond to the caches used to evaluate client membership in a
netgroup. To show the netgroup cache, you must specify the following item:

* Vserver: The name of the Vserver whose netgroup cache you want to
display.

The following information is displayed per cache entry:

* Vserver name: The name of the Vserver.
* Netgroup name: The name of the netgroup.
* State of the cache entry: The state of the cache entry. There are four
possible values:

  * initializing: The cache entry is being populated for the first time.
  * ready: Processing of the cache entry is complete and it is ready to be
used.
  * not-found: The netgroup could not be found.
  * abandoned: The cache entry has been abandoned.


* Total number of hosts in the netgroup cache: The number of host names
retrieved from the name service in mapping the netgroup to a list of
hosts.
* How long it took to expand the netgroup: How long it took to expand the
netgroup the last time in the queue.
* Entry is refreshing: If the entry is a complete miss or refresh.
* Next refresh time: When the next refresh is scheduled to take place.
* Netgroup by host state: Boolean state indicating if netgroup-by-host
feature is used for resolving netgroup membership check.
* Number of IP addresses cached: Number of client IP addresses that are
matched for the netgroup. The count includes both positive and negative
results.



== 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.
```

```
``-vserver <vserver name>`` - Vserver::
If you specify this parameter, the command displays the netgroup cache
information only if the Vserver name matches the specified value.

``[-netgroup <text>]`` - Name of the Netgroup::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup name matches the specified value.

``[-cache-state {initializing|ready|not-found|abandoned}]`` - State of the
Cache Entry::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup cache state matches the specified value.

``[-total-hosts <integer>]`` - Total Number of Hosts in the Netgroup::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup record's count of host names matches the
specified value.

``[-expansion-duration <[[<hours>:]<minutes>:]<seconds>>]`` - Expansion
Duration::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup record expansion time matches the
specified value.

``[-is-refreshing {true|false}]`` - Is Entry Refreshing?::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup record refreshing state matches the
specified value.

``[-time-next-refresh <Date>]`` - Next Refresh Time::
If you specify this parameter, the command displays the netgroup cache
information only if the time of the next scheduled refresh matches the
specified value.

``[-num-ip-addrs-cache <integer>]`` - Number of Cached IP Addresses::
If you specify this parameter, the command displays the netgroup cache
information only if the number of cached IP addresses matches the
specified value.

== Examples

The following example displays the netgroup cache for the Vserver vs1 and
the netgroup netgroup1:
```

cluster1::> vserver export-policy netgroup cache show -vserver vs1 -netgroup netgroup1
Vserver Netgroup State

-------- --------- ------------
vs1 netgroup1 Ready

```
[[ID3cee017c01d26a5f75f85fb53aaba08d]]
= vserver export-policy netgroup queue show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the Netgroup Processing Queue

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver export-policy netgroup queue show` command displays the
ongoing processing of the netgroup cache for a node. Entries shown here
are not used to evaluate client membership in a netgroup. The following
information is displayed per queue entry:

* Vserver name: The name of the Vserver.
* Netgroup name: The name of the netgroup.
* Age of entry in the queue: How long the entry has been in the queue.
* Queue state: The state of the entry in the queue. There are three
possible values:

  * running: The entry is currently being processed.
  * waiting: The entry is waiting to be processed.
  * retrying: The entry is waiting to be reprocessed.

 Note that as the `vserver export-policy netgroup queue show` command is
not atomic. Several queue entries might show up in the 'running' state.
* Number of times retried in the queue: The number of times was the entry
was taken off of the netgroup processing queue and added back on it.
* Total number of hosts in the netgroup: The number of host names
retrieved from the name service in mapping the netgroup to a list of
hosts.



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays the netgroup cache
information only if the Vserver name matches the specified value.

``[-netgroup <text>]`` - Name of the Netgroup::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup name matches the specified value.

``[-queue-state {waiting|running|retrying}]`` - State of Entry in the
Queue::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup queue state matches the specified value.

``[-age <[[<hours>:]<minutes>:]<seconds>>]`` - Age of Entry in the Queue::
If you specify this parameter, the command displays the netgroup cache
information only if the age of when the netgroup record was put on the
netgroup processing queue matches the specified value.

``[-retries-on-queue <integer>]`` - Number of Retries on the Queue::
If you specify this parameter, the command displays the netgroup cache
information only if, during a refresh, the number of times the netgroup
record has been put back on the netgroup processing queue matches the
specified value.

``[-total-hosts <integer>]`` - Total Number of Hosts in the Netgroup::
If you specify this parameter, the command displays the netgroup cache
information only if the netgroup record's count of hosts matches the
specified value.

== Examples

The following example displays the netgroup queue:
```

cluster1::> vserver export-policy netgroup queue show
Age on Total
Vserver Netgroup State Queue Hosts
--- --- --- --- ---
testvs1 test-netgr retrying 0:0:47 12441
testvs1 test waiting 0:01:35 -

```
[[ID2389d6b92609e30ee9381f309e2d221e]]
= vserver export-policy rule add-clientmatches
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add list of clientmatch strings to an existing rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rule add-clientmatches` command adds a list of
strings to the clientmatch field of a specified export rule in a policy.
This command only operates on the clientmatch field; to modify other
fields in a rule use the `vserver export-policy modify` command.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy is
located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the name of the export policy containing the
export rule to which you want to add additional clientmatch strings.

``-ruleindex <integer>`` - Rule Index::
This parameter specifies the index number of the export rule to which you
want to add additional clientmatch strings. To view a list of rules with
their index numbers, use the xref:{relative_path}vserver-export-policy-
rule-show.html[vserver export-policy rule show] command.

``-clientmatches <text>`` - List of Clientmatch Strings to Add::
This parameter specifies list of the match strings specifying the client
or clients to add to the export rule. Duplicate match strings will not be
created and the list may not contain duplicates entries. Match strings
from the clientmatches list are added to the clientmatch field if the
match string is not identical to one of the strings already in the
clientmatch field. The maximum number of clientmatches that can be created
is 4096. You can specify the match string in any of the following formats:
+
  * As a hostname; for instance, host1
```

```
  * As an IPv4 address; for instance, 10.1.12.24
  * As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
  * As an IPv4 address with a subnet mask expressed as a number of bits;
for instance, 10.1.12.0/24
  * As an IPv6 address with a subnet mask expressed as a number of bits;
for instance, fd20:8b1e:b255:4071::/64
  * As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
  * As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
  * As a domain name preceded by the . character; for instance,
.example.com

+
Note: Entering an IP address range, such as 10.1.12.10-10.1.12.70, is not
allowed. Entries in this format are interpreted as a text string and
treated as a hostname.

== Examples

The following example adds match strings "2.2.2.2" and "3.3.3.3" to the
clientmatch field of the export rule with index number 3 in an export
policy named default_expolicy on a Vserver named vs0.
```

cluster1::> vserver export-policy rule add-clientmatches -vserver vs0 -policyname default_expolicy -ruleindex 3 -clientmatches "2.2.2.2,3.3.3.3"

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-show.html[vserver export-
policy rule show]




[[ID29066ed8cb2f089b3f257f17c50b95ea]]
= vserver export-policy rule create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a rule

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rule create` command creates an export rule and
adds it to a policy. To create an export rule, you must specify the
following items:

* Vserver
* Export policy
* Clients that match the rule
* Read-only access rule
* Read-write access rule

You can optionally specify the following items:

* Index number; that is, the location of the export rule in the policy
* Access protocol
* Anonymous ID
* Superuser security type
* Whether suid access is enabled
* Whether creation of devices is enabled
* Whether UNIX-type permissions changes on NTFS (Windows) volumes are
prohibited or allowed when the request originates from an NFS client
(advanced privilege and higher only)
* Whether ownership changes are restricted or not (advanced privilege and
higher only)



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy is
located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the name of the export policy to which you want
to add the new export rule. The export policy must already exist. To
create an export policy, see the xref:{relative_path}vserver-export-
policy-create.html[vserver export-policy create] command.

``[-ruleindex <integer>]`` - Rule Index::
This optional parameter specifies the index number of the export rule that
you want to create. If you specify an index number that already matches a
rule, the index number of the existing rule is incremented, as are the
index numbers of all subsequent rules, either to the end of the list or to
```

```
an open space in the list. If you do not specify an index number, the new
rule is placed at the end of the policy's list. Valid values are values
from 1 to 4294967295.

``[-protocol <Client Access Protocol>,...]`` - Access Protocol::
This optional parameter specifies the list of access protocols for which
you want to apply the export rule. Possible values include the following:
+
  * ``_any_`` - Any current or future access protocol
  * ``_nfs_`` - Any current or future version of NFS
  * ``_nfs3_`` - The NFSv3 protocol
  * ``_nfs4_`` - The NFSv4 protocol
  * ``_cifs_`` - The CIFS protocol

+
You can specify a comma-separated list of multiple access protocols for an
export rule. If you specify the protocol as any, you cannot specify any
other protocols in the list. If you do not specify this parameter, the
value defaults to ``_any_`` . If you enable NFSv4, you will not be able to
apply the policy to which this rule belongs to a FlexGroup, as FlexGroup
volumes do not support NFSv4 protocol access.

``-clientmatch <text>`` - List of Client Match Hostnames, IP Addresses,
Netgroups, or Domains::
This parameter specifies list of the match strings specifying the client
or clients to which the export rule applies. Duplicate match strings in
the same rule are not allowed. The maximum number of clientmatches that
can be created is 4096. You can specify the match string in any of the
following formats:
+
  * As a hostname; for instance, host1
  * As an IPv4 address; for instance, 10.1.12.24
  * As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
  * As an IPv4 address with a subnet mask expressed as a number of bits;
for instance, 10.1.12.0/24
  * As an IPv6 address with a subnet mask expressed as a number of bits;
for instance, fd20:8b1e:b255:4071::/64
  * As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
  * As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
  * As a domain name preceded by the . character; for instance,
.example.com

+
Note: Entering an IP address range, such as 10.1.12.10-10.1.12.70, is not
```

```
allowed. Entries in this format are interpreted as a text string and
treated as a hostname.

``-rorule <authentication method>,...`` - RO Access Rule::
This parameter specifies the security type for read-only access to volumes
that use the export rule. Possible values include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is AUTH_SYS. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes sys.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5. The effective security type of
the incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes krb5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with integrity service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5i.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with privacy service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5p.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is CIFS NTLM. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes ntlm.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume regardless of the
security type of that incoming request. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) remains the same as the security type of the incoming
request.
+
NOTE: If the security type of the incoming request is AUTH_NONE, read
access will be granted to that incoming request as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume as an anonymous user
if the security type of that incoming request is not explicitly listed in
```

```
the list of values in the rorule. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes none.
  * ``_never_`` - For an incoming request from a client matching the
clientmatch criteria, do not allow any access to the volume regardless of
the security type of that incoming request.

+
You can specify a comma-separated list of multiple security types for an
export rule. If you specify the security type as ``_any_`` or ``_never_``
, you cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria, if the security type doesn't match any of the values listed in
rorule (as explained above), access will be denied to that incoming
request.
+



``-rwrule <authentication method>,...`` - RW Access Rule::
This parameter specifies the security type for read-write access to
volumes that use the export rule. Possible values include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
AUTH_SYS.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
Kerberos v5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
Kerberos v5 with integrity service.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
Kerberos v5 with privacy service.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is CIFS
NTLM.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume regardless of the
effective security type (determined from rorule) of that incoming request.
```

```
+
NOTE: If the effective security type (determined from rorule) of the
incoming request is none, write access will be granted to that incoming
request as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume as an anonymous
user if the effective security type (determined from rorule) of that
incoming request is none.
  * ``_never_`` - For an incoming request from a client matching the
clientmatch criteria, do not allow write access to the volume regardless
of the effective security type (determined from rorule) of that incoming
request.

+
You can specify a comma-separated list of multiple security types for an
export rule. If you specify the security type as ``_any_`` or ``_never_``
, you cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria, if the effective security type (determined by rorule) doesn't
match any of the values listed in rwrule (as explained above), write
access will be denied to that incoming request.
+


``[-anon <text>]`` - User ID To Which Anonymous Users Are Mapped::
This parameter specifies a UNIX user ID or user name that the user
credentials are mapped to when evaluation of rorule or superuser
parameters result in user being mapped to the anonymous user. The default
setting of this parameter is 65534, which is normally associated with the
user name "nobody" or "nfsnobody" in Linux environments. NetApp appliances
use 65534 as the user "pcuser", which is generally used for multiprotocol
operations. Because of this difference, if using local files and NFSv4,
the name string for users mapped to 65534 might not match. This
discrepancy might cause files to be written as the user specified in the
/etc/idmapd.conf file on the client (Linux) or /etc/default/nfs file
(Solaris), particularly when using multiprotocol (CIFS and NFS) on the
same datasets. The following notes apply to the use of this parameter:
+
  * To disable access by any client with a user ID of 0, specify a value
of 65535 which is associated with the user nobody.
```

```
``[-superuser <authentication method>,...]`` - Superuser Security Types::
This parameter specifies a security type for superuser access to files.
The default setting of this parameter is ``_none_`` . Possible values
include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is AUTH_SYS.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is Kerberos v5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is Kerberos v5 with integrity service.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is Kerberos v5 with privacy service.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is CIFS NTLM.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume regardless of the effective security type (determined by rorule) of
that incoming request.
+
NOTE: If the effective security type (determined from rorule) of the
incoming request is none, access will be granted to that incoming request
as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow access to the volume as
an anonymous user if the effective security type (determined from rorule)
of that incoming request is none.

+
You can specify a comma-separated list of multiple security types for
superuser access. If you specify the security type as ``_any_`` , you
cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
```

```
criteria and with the user ID 0, if the effective security type doesn't
match any of the values listed in superuser (as explained above), the user
ID is mapped to anonymous user.
+


``[-allow-suid {true|false}]`` - Honor SetUID Bits in SETATTR::
This parameter specifies whether set user ID (suid) and set group ID
(sgid) access is enabled by the export rule. The default setting is `true`
.

``[-allow-dev {true|false}]`` - Allow Creation of Devices::
This parameter specifies whether the creation of devices is enabled by the
export rule. The default setting is `true` .

``[-ntfs-unix-security-ops {ignore|fail}]`` - NTFS Unix Security Options
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether UNIX-type permissions changes on NTFS
(Windows) volumes are prohibited (fail) or allowed (ignore) when the
request originates from an NFS client. The default setting is `fail` .

``[-chown-mode {restricted|unrestricted}]`` - Change Ownership Mode
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies who is allowed to change the ownership mode of a
file. The default setting is `restricted` . The allowed values are:
+
  * restricted - Only root may change the ownership of the file.
  * unrestricted - Non-root users may change file ownership provided the
on-disk permissions allow the operation.



== Examples

The following example creates an export rule with index number 1 in an
export policy named read_only_expolicy on a Vserver named vs0. The rule
matches all clients in the domains named example.com or example.net. The
rule enables all access protocols. It enables read-only access by any
matching client and requires authentication by AUTH_SYS, NTLM, or Kerberos
5 for read-write access. Clients with the UNIX user ID zero are mapped to
user ID 65534 (which normally maps to the user name nobody). It does not
enable suid and sgid access or the creation of devices.
```

cluster1::> vserver export-policy rule create -vserver vs0 -policyname read_only_expolicy -ruleindex 1 -protocol any -clientmatch ".example.com,.example.net" -rorule any -rwrule "ntlm,krb5,sys" -anon 65534 -allow -suid false

-allow-dev false

```
== Related Links

* xref:{relative_path}vserver-export-policy-create.html[vserver export-
policy create]




[[IDc7448033ec73200a4b55888fb84f0c35]]
= vserver export-policy rule delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rule delete` command deletes an export rule
from a policy. You can specify the export rule by specifying its index
number in the policy. When you delete a rule, the other rules in the
policy are not automatically renumbered or reordered. You can use the
xref:{relative_path}vserver-export-policy-rule-setindex.html[vserver
export-policy rule setindex] command to reorder the rules in a rule set.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver which contains the export policy.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the export policy from which you want to delete a
rule.

``-ruleindex <integer>`` - Rule Index::
This parameter specifies the index number of the rule that you want to
delete. You can use the xref:{relative_path}vserver-export-policy-rule-
show.html[vserver export-policy rule show] command to view a list of rules
with their index numbers.
```

```
== Examples

The following example deletes an export rule with the index number 5 from
an export policy named rs1 on a Vserver named vs0:
```

cluster1::> vserver export-policy rule delete -vserver vs0
-policyname read_only_expolicy -ruleindex 5

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-setindex.html[vserver
export-policy rule setindex]
* xref:{relative_path}vserver-export-policy-rule-show.html[vserver export-
policy rule show]




[[IDc5c9c62605aa29df10afd9e072817157]]
= vserver export-policy rule modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rule modify` command modifies a specified
export rule in a policy. This command cannot change the position of a rule
in a policy; to reorder rules in a policy, use the
xref:{relative_path}vserver-export-policy-rule-setindex.html[vserver
export-policy rule setindex] command. Duplicate match strings in the same
rule are not allowed. You can use this command to change the following
attributes of an export rule:

* Access protocol
* Client match specification
* Read-only access rule
* Read-write access rule
```

```
* Anonymous ID
* Superuser security type
* Whether suid access is enabled
* Whether creation of devices is enabled
* Whether UNIX-type permissions changes on NTFS (Windows) volumes are
prohibited or allowed when the request originates from an NFS client
(advanced privilege and higher only)
* Whether ownership changes are restricted or not (advanced privilege and
higher only)



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy is
located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the name of the export policy containing the
export rule that you want to modify.

``-ruleindex <integer>`` - Rule Index::
This parameter specifies the index number of the export rule that you want
to modify. To view a list of rules with their index numbers, use the
xref:{relative_path}vserver-export-policy-rule-show.html[vserver export-
policy rule show] command.

``[-protocol <Client Access Protocol>,...]`` - Access Protocol::
This optional parameter specifies the list of access protocols for which
you want to apply the export rule. Possible values include the following:
+
  * ``_any_`` - Any current or future access protocol
  * ``_nfs_`` - Any current or future version of NFS
  * ``_nfs3_`` - The NFSv3 protocol
  * ``_nfs4_`` - The NFSv4 protocol
  * ``_cifs_`` - The CIFS protocol

+
You can specify a comma-separated list of multiple access protocols for an
export rule. If you specify the protocol as any, you cannot specify any
other protocols in the list. If you do not specify this parameter, the
value defaults to ``_any_`` . If you enable NFSv4, you will not be able to
apply the policy to which this rule belongs to a FlexGroup, as FlexGroup
volumes do not support NFSv4 protocol access.
```

```
``[-clientmatch <text>]`` - List of Client Match Hostnames, IP Addresses,
Netgroups, or Domains::
This parameter specifies list of the match strings specifying the client
or clients to which the export rule applies. The maximum number of
clientmatches that can be created is 4096. You can specify the match
string in any of the following formats:
+
  * As a hostname; for instance, host1
  * As an IPv4 address; for instance, 10.1.12.24
  * As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
  * As an IPv4 address with a subnet mask expressed as a number of bits;
for instance, 10.1.12.0/24
  * As an IPv6 address with a subnet mask expressed as a number of bits;
for instance, fd20:8b1e:b255:4071::/64
  * As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
  * As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
  * As a domain name preceded by the . character; for instance,
.example.com

+
Note: Entering an IP address range, such as 10.1.12.10-10.1.12.70, is not
allowed. Entries in this format are interpreted as a text string and
treated as a hostname.

``[-rorule <authentication method>,...]`` - RO Access Rule::
This parameter modifies the security type for read-only access to volumes
that use the export rule. Possible values include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is AUTH_SYS. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes sys.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5. The effective security type of
the incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes krb5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with integrity service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5i.
  * ``_krb5p_`` - For an incoming request from a client matching the
```

```
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with privacy service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5p.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is CIFS NTLM. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes ntlm.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume regardless of the
security type of that incoming request. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) remains the same as the security type of the incoming
request.
+
NOTE: If the security type of the incoming request is AUTH_NONE, read
access will be granted to that incoming request as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume as an anonymous user
if the security type of that incoming request is not explicitly listed in
the list of values in the rorule. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes none.
  * ``_never_`` - For an incoming request from a client matching the
clientmatch criteria, do not allow any access to the volume regardless of
the security type of that incoming request.

+
You can specify a comma-separated list of multiple security types for an
export rule. If you specify the security type as ``_any_`` or ``_never_``
, you cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria, if the security type doesn't match any of the values listed in
rorule (as explained above), access will be denied to that incoming
request.
+



``[-rwrule <authentication method>,...]`` - RW Access Rule::
This parameter modifies the security type for read-write access to volumes
that use the export rule. Possible values include the following:
+
```

```
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
AUTH_SYS.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
Kerberos v5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
Kerberos v5 with integrity service.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
Kerberos v5 with privacy service.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is CIFS
NTLM.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume regardless of the
effective security type (determined from rorule) of that incoming request.
+
NOTE: If the effective security type (determined from rorule) of the
incoming request is none, write access will be granted to that incoming
request as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume as an anonymous
user if the effective security type (determined from rorule) of that
incoming request is none.
  * ``_never_`` - For an incoming request from a client matching the
clientmatch criteria, do not allow write access to the volume regardless
of the effective security type (determined from rorule) of that incoming
request.

+
You can specify a comma-separated list of multiple security types for an
export rule. If you specify the security type as ``_any_`` or ``_never_``
, you cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria, if the effective security type (determined by rorule) doesn't
match any of the values listed in rwrule (as explained above), write
```

```
access will be denied to that incoming request.
+


``[-anon <text>]`` - User ID To Which Anonymous Users Are Mapped::
This parameter specifies a UNIX user ID or user name that the user
credentials are mapped to when evaluation of rorule or superuser
parameters result in user being mapped to the anonymous user. The default
setting of this parameter is 65534, which is normally associated with the
user name "nobody" or "nfsnobody" in Linux environments. NetApp appliances
use 65534 as the user "pcuser", which is generally used for multiprotocol
operations. Because of this difference, if using local files and NFSv4,
the name string for users mapped to 65534 might not match. This
discrepancy might cause files to be written as the user specified in the
/etc/idmapd.conf file on the client (Linux) or /etc/default/nfs file
(Solaris), particularly when using multiprotocol (CIFS and NFS) on the
same datasets. The following notes apply to the use of this parameter:
+
  * To disable access by any client with a user ID of 0, specify a value
of 65535 which is associated with the user nobody.



``[-superuser <authentication method>,...]`` - Superuser Security Types::
This parameter specifies a security type for superuser access to files.
The default setting of this parameter is ``_none_`` . Possible values
include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is AUTH_SYS.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is Kerberos v5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is Kerberos v5 with integrity service.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is Kerberos v5 with privacy service.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
```

```
volume if the effective security type (determined from rorule) of that
incoming request is CIFS NTLM.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume regardless of the effective security type (determined by rorule) of
that incoming request.
+
NOTE: If the effective security type (determined from rorule) of the
incoming request is none, access will be granted to that incoming request
as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow access to the volume as
an anonymous user if the effective security type (determined from rorule)
of that incoming request is none.

+
You can specify a comma-separated list of multiple security types for
superuser access. If you specify the security type as ``_any_`` , you
cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria and with the user ID 0, if the effective security type doesn't
match any of the values listed in superuser (as explained above), the user
ID is mapped to anonymous user.
+


``[-allow-suid {true|false}]`` - Honor SetUID Bits in SETATTR::
This parameter specifies whether set user ID (suid) and set group ID
(sgid) access is enabled by the export rule. The default setting is `true`
.

``[-allow-dev {true|false}]`` - Allow Creation of Devices::
This parameter specifies whether the creation of devices is enabled by the
export rule. The default setting is `true` .

``[-ntfs-unix-security-ops {ignore|fail}]`` - NTFS Unix Security Options
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether UNIX-type permissions changes on NTFS
(Windows) volumes are prohibited (with value ``_fail_`` ) or allowed (with
value ``_ignore_`` ) when the request originates from an NFS client. The
default setting is `fail` . This parameter is only used if you set the
NTFS UNIX security option for the Vserver to `use_export_policy` ;
otherwise, it has no effect.
```

```
``[-chown-mode {restricted|unrestricted}]`` - Change Ownership Mode
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies who is authorized to change the ownership mode of
a file. The default setting is `restricted` . This parameter is only used
if you set the change ownership mode option for the Vserver to
`use_export_policy` ; otherwise, it has no effect. The allowed values are
:
+
  * restricted - Only root user can change the ownership of the file.
  * unrestricted - Non-root users may change file ownership provided the
on-disk permissions allow the operation.



== Examples

The following example modifies the export rule with index number 3 in an
export policy named default_expolicy on a Vserver named vs0. The rule is
modified to match any clients in the netgroups named group1 or group2 to
enable NFSv2 and CIFS support, to enable read-only access by any matching
client, to require authentication by NTLM or Kerberos 5 for read-write
access, and to enable suid and sgid access.
```

cluster1::> vserver export-policy rule modify -vserver vs0 -policyname default_expolicy -ruleindex 3 -protocol "nfs2,cifs"
-clientmatch "@group1, @group2" -rorule any -rwrule "ntlm,krb5" -allow-suid true

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-setindex.html[vserver
export-policy rule setindex]
* xref:{relative_path}vserver-export-policy-rule-show.html[vserver export-
policy rule show]




[[IDc60cfba99dacf0f8b1b60c10e7f0e120]]
= vserver export-policy rule remove-clientmatches
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove list of clientmatch strings from an existing rule
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rule remove-clientmatches` command removes a
list of strings from the clientmatch field of a specified export rule in a
policy. This command only operates on the clientmatch field; to modify
other fields in a rule use the `vserver export-policy modify` command.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy is
located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the name of the export policy containing the
export rule from which you want to remove clientmatch strings.

``-ruleindex <integer>`` - Rule Index::
This parameter specifies the index number of the export rule from which
you want to remove clientmatch strings. To view a list of rules with their
index numbers, use the xref:{relative_path}vserver-export-policy-rule-
show.html[vserver export-policy rule show] command.

``-clientmatches <text>`` - List of Clientmatch Strings to Remove::
This parameter specifies list of the match strings specifying the client
or clients to remove from the export rule. Match strings are removed from
the clientmatch field if the match string is identical to one of the
elements in the clientmatches list. If all match strings are removed from
the clientmatch field the entire export rule is deleted. You can specify
the match string in any of the following formats:
+
  * As a hostname; for instance, host1
  * As an IPv4 address; for instance, 10.1.12.24
  * As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
  * As an IPv4 address with a subnet mask expressed as a number of bits;
for instance, 10.1.12.0/24
  * As an IPv6 address with a subnet mask expressed as a number of bits;
for instance, fd20:8b1e:b255:4071::/64
  * As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
  * As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
```

```
  * As a domain name preceded by the . character; for instance,
.example.com

+
Note: Entering an IP address range, such as 10.1.12.10-10.1.12.70, is not
allowed. Entries in this format are interpreted as a text string and
treated as a hostname.

== Examples

The following example removes match strings "2.2.2.2" and "3.3.3.3" from
the clientmatch field of the export rule with index number 3 in an export
policy named default_expolicy on a Vserver named vs0.
```

cluster1::> vserver export-policy rule remove-clientmatches -vserver vs0 -policyname default_expolicy -ruleindex 3 -clientmatches "2.2.2.2,3.3.3.3"

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-show.html[vserver export-
policy rule show]




[[ID848d2fac91023d45fdbf925b7362b674]]
= vserver export-policy rule setindex
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Move a rule to a specified index

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rule setindex` command modifies the index
number of the specified export rule. If the new index number is already in
use, the command reorders the list to accommodate it. If the existing
index is given a higher index number (that is, later in the list), the
command decrements the index numbers of rules between the moved rule and
moved-to rule; otherwise, the command increments the index numbers between
```

```
the moved-to rule and the existing rule.

You can use the xref:{relative_path}vserver-export-policy-rule-
show.html[vserver export-policy rule show] command to view a list of rules
with their index numbers.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the export policy is
located.

``-policyname <export policy name>`` - Policy Name::
This parameter specifies the export policy that contains the rule whose
index number you want to modify.

``-ruleindex <integer>`` - Rule Index::
This parameter specifies the index number of the rule that you want to
move.

``-newruleindex <integer>`` - Index::
This parameter specifies the new index number for the rule.

== Examples

The following example changes the index number of a rule at index number 5
to index number 3 in an export policy named rs1 on a Vserver named vs0:
```

cluster1::> vserver export-policy rule setindex -vserver vs0
-policyname read_only_policy -ruleindex 5 -newruleindex 3

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-show.html[vserver export-
policy rule show]




[[ID0691d4931af32fc2803b2a73b966a657]]
= vserver export-policy rule show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display a list of rules

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver export-policy rule show` command displays information about
export rules. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information:

* Vserver name
* Export policy name
* Export rule index number
* Access protocol
* Client match
* Read-only access rule
* Read-write access rule

To display detailed information about a specific export rule, run the
command with the `-vserver` , `-policyname` , and `-ruleindex` parameters.
The detailed view provides all of the information in the previous list and
the following additional information:

* Anonymous ID
* Superuser security type
* Whether set user ID (suid) and set group ID (sgid) access is enabled
* Whether creation of devices is enabled
* NTFS security settings
* Change ownership mode

You can specify additional parameters to display only the information that
matches those parameters. For example, to display information only about
export rules that have a read-write rule value of never, run the command
with the `-rwrule never` parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.
```

```
``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the `-policyname` parameter, and the `-
ruleindex` parameter, the command displays detailed information about the
specified export rule. If you specify this parameter by itself, the
command displays information only about the export rules on the specified
Vserver.

``[-policyname <export policy name>]`` - Policy Name::
If you specify this parameter, the `-vserver` parameter, and the `-
ruleindex` parameter, the command displays detailed information about the
specified export rule. If you specify this parameter by itself, the
command displays information only about the export rules on the specified
policy.

``[-ruleindex <integer>]`` - Rule Index::
If you specify this parameter, the `-vserver` parameter, and the `-
policyname` parameter, the command displays detailed information about the
specified export rule. If you specify this parameter by itself, the
command displays information only about the export rules that have the
specified index number.

``[-protocol <Client Access Protocol>,...]`` - Access Protocol::
If you specify this parameter, the command displays information only about
the export rules that have the specified access protocol or protocols.
Possible values include the following:
+
  * ``_any_`` - Any current or future access protocol
  * ``_nfs_`` - Any current or future version of NFS
  * ``_nfs3_`` - The NFSv3 protocol
  * ``_nfs4_`` - The NFSv4 protocol
  * ``_cifs_`` - The CIFS protocol

+
You can specify a comma-separated list of multiple access protocols for an
export rule. If you specify the protocol as any, you cannot specify any
other protocols in the list.

``[-clientmatch <text>]`` - List of Client Match Hostnames, IP Addresses,
Netgroups, or Domains::
If you specify this parameter, the command displays information only about
the export rules that have a clientmatch list containing all of the
strings in the specified client match. You can specify the match as a list
of strings in any of the following formats:
+
  * As a hostname; for instance, host1
```

```
  * As an IPv4 address; for instance, 10.1.12.24
  * As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
  * As an IPv4 address with a subnet mask expressed as a number of bits;
for instance, 10.1.12.0/24
  * As an IPv6 address with a subnet mask expressed as a number of bits;
for instance, fd20:8b1e:b255:4071::/64
  * As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
  * As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
  * As a domain name preceded by the . character; for instance,
.example.com



``[-rorule <authentication method>,...]`` - RO Access Rule::
If you specify this parameter, the command displays information only about
the export rule or rules that have the specified read-only rule. Possible
values include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is AUTH_SYS. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes sys.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5. The effective security type of
the incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes krb5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with integrity service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5i.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with privacy service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5p.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is CIFS NTLM. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes ntlm.
  * ``_any_`` - For an incoming request from a client matching the
```

```
clientmatch criteria, allow read access to the volume regardless of the
security type of that incoming request. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) remains the same as the security type of the incoming
request.
+
NOTE: If the security type of the incoming request is AUTH_NONE, read
access will be granted to that incoming request as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume as an anonymous user
if the security type of that incoming request is not explicitly listed in
the list of values in the rorule. The effective security type of the
incoming request (to be used subsequently in evaluation of
rwrule/superuser) becomes none.
  * ``_never_`` - For an incoming request from a client matching the
clientmatch criteria, do not allow any access to the volume regardless of
the security type of that incoming request.

+
You can specify a comma-separated list of multiple security types for an
export rule. If you specify the security type as ``_any_`` or ``_never_``
, you cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria, if the security type doesn't match any of the values listed in
rorule (as explained above), access will be denied to that incoming
request.
+


``[-rwrule <authentication method>,...]`` - RW Access Rule::
If you specify this parameter, the command displays information only about
the export rule or rules that have the specified read-write rule. Possible
values include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
AUTH_SYS.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is
Kerberos 5.
  * ``_krb5i_`` - For an incoming request from a client matching the
```

```
clientmatch criteria, allow write access to the volume if the security
type of that incoming request is Kerberos v5 with integrity service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5i.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the security
type of that incoming request is Kerberos v5 with privacy service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5p.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume if the effective
security type (determined from rorule) of that incoming request is CIFS
NTLM.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume regardless of the
effective security type (determined from rorule) of that incoming request.
+
NOTE: If the effective security type (determined from rorule) of the
incoming request is none, write access will be granted to that incoming
request as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria, allow write access to the volume as an anonymous
user if the effective security type (determined from rorule) of that
incoming request is none.
  * ``_never_`` - For an incoming request from a client matching the
clientmatch criteria, do not allow write access to the volume regardless
of the effective security type (determined from rorule) of that incoming
request.

+
You can specify a comma-separated list of multiple security types for an
export rule. If you specify the security type as ``_any_`` or ``_never_``
, you cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria, if the effective security type (determined by rorule) doesn't
match any of the values listed in rwrule (as explained above), write
access will be denied to that incoming request.
+


``[-anon <text>]`` - User ID To Which Anonymous Users Are Mapped::
If you specify this parameter, the command displays information only about
the export rule or rules that have the specified anonymous ID.
```

```
``[-superuser <authentication method>,...]`` - Superuser Security Types::
If you specify this parameter, the command displays information only about
the export rule or rules that have the specified superuser security type.
Possible values include the following:
+
  * ``_sys_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is AUTH_SYS.
  * ``_krb5_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is Kerberos v5.
  * ``_krb5i_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with integrity service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5i.
  * ``_krb5p_`` - For an incoming request from a client matching the
clientmatch criteria, allow read access to the volume if the security type
of that incoming request is Kerberos v5 with privacy service. The
effective security type of the incoming request (to be used subsequently
in evaluation of rwrule/superuser) becomes krb5p.
  * ``_ntlm_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume if the effective security type (determined from rorule) of that
incoming request is CIFS NTLM.
  * ``_any_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow superuser access to the
volume regardless of the effective security type (determined by rorule) of
that incoming request.
+
NOTE: If the effective security type (determined from rorule) of the
incoming request is none, access will be granted to that incoming request
as an anonymous user.
+

  * ``_none_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow access to the volume as
an anonymous user if the effective security type (determined from rorule)
of that incoming request is none.
  * ``_never_`` - For an incoming request from a client matching the
clientmatch criteria and with the user ID 0, allow access to the volume as
an anonymous user regardless of the effective security type (determined
from rorule) of that incoming request.
+
```

```
NOTE: Only export rules that were created in an earlier release can have
the superuser parameter set to the security type ``_never_``
+



+
You can specify a comma-separated list of multiple security types for
superuser access. If you specify the security type as ``_any_`` , you
cannot specify any other security types.
+
NOTE: For an incoming request from a client matching the clientmatch
criteria and with the user ID 0, if the effective security type doesn't
match any of the values listed in superuser (as explained above), the user
ID is mapped to anonymous user.
+



``[-allow-suid {true|false}]`` - Honor SetUID Bits in SETATTR::
If you specify this parameter, the command displays information only about
the export rule or rules that have the specified setting for set user ID
(suid) and set group ID (sgid) access.

``[-allow-dev {true|false}]`` - Allow Creation of Devices::
If you specify this parameter, the command displays information only about
the export rule or rules that have the specified setting for the creation
of devices.

``[-ntfs-unix-security-ops {ignore|fail}]`` - NTFS Unix Security Options
[.advaced-priv]#(privilege: advanced)#::
If you have specified this parameter for a particular export policy rule,
then the command displays information about the UNIX security options that
apply to that export policy rule. The setting can either prohibit (with
value ``_fail_`` ) or allow (with value ``_ignore_`` ) UNIX-type
permissions changes on NTFS (Windows) volumes when the request originates
from an NFS client. If the Vserver NTFS UNIX security option is set to
fail or allow for the Vserver, then this parameter is overridden.

``[-ntfs-unix-security-ops-vs {fail|ignore|use_export_policy}]`` - Vserver
NTFS Unix Security Options [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
UNIX security options that apply to all volumes in this Vserver. The
setting can prohibit (with value ``_fail_`` ) or allow (with value
``_ignore_`` ) UNIX-type permissions changes on NTFS (Windows) volumes
when the request originates from an NFS client, or you can set it to
`use_export_policy` . If you set this parameter to `fail` or `allow` ,
this parameter overrides the individual UNIX security options set for the
```

```
export policy rules. If you set this parameter to `use_export_policy` ,
the UNIX security options associated with the respective export policy
rule is used.

``[-chown-mode {restricted|unrestricted}]`` - Change Ownership Mode
[.advaced-priv]#(privilege: advanced)#::
If you have specified this parameter for a particular export policy rule,
then the command displays information about the change ownership mode that
applies to that export-policy rule. The setting can either allow only the
root (with value ``_restricted_`` ) or all users (with value
``_unrestricted_`` ) to change file ownership provided the on-disk
permissions allow the operation. If the Vserver change ownership mode is
set to restricted or unrestricted for the Vserver, then this parameter is
overridden.

``[-chown-mode-vs {restricted|unrestricted|use_export_policy}]`` - Vserver
Change Ownership Mode [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
change ownership mode that applies to all volumes in this Vserver. The
setting can allow only the root (with value ``_restricted_`` ) or all
users (with value ``_unrestricted_`` ) to change ownership of the files
that they own, or you can set it to `use_export_policy` . If you set this
parameter to `restricted` or `unrestricted` , this parameter overrides the
individual change ownership mode set for the export policy rules. If you
set this parameter to `use_export_policy` , the change ownership mode
associated with the respective export policy rule is used.

== Examples

The following example displays information about all export rules:


cluster1::> vserver export-policy rule show
Policy Rule Access Client RO
Vserver Name Index Protocol Match Rule
------------ ------------------ ------ -------- ------------------------ ------
vs0 default_expolicy 1 any 0.0.0.0/0,::0/0 any
vs0 read_only_expolicy 2 any 0.0.0.0/0 any
vs1 default_expolicy 1 any 10.10.10.10,11.11.11.11 any
vs1 test_expolicy 1 any 0.0.0.0/0 any
4 entries were displayed.
```

```
:leveloffset: -1



= vserver fcp commands

:leveloffset: +1
```

```
[[ID01d46d7ca5a8085e8ba9901f86102ebd]]
= vserver fcp create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create FCP service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates an FCP service for a Vserver. An FCP service must be
licensed before you can manage FCP services. If the FCP service is not
licensed, the FCP command returns an error.

When you create an FCP service on a Vserver, the Vserver has the following
configuration defaults:

* The administrative status of the FCP service is ``_up_`` .
* The FCP command automatically generates a unique World Wide Node Name
(WWNN).



== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the FCP service.

``[-target-name <text>]`` - Target Name [.advaced-priv]#(privilege:
advanced)#::
The FCP World Wide Node Name (WWNN) for the service. All FCP LIFs in the
Vserver will share the specified WWNN. The format for a WWNN is
"XX:XX:XX:XX:XX:XX:XX:XX" where X is a hexadecimal digit.
+
Unless the `force` option is also provided, the specified WWNN must be
within one of the following vendor registered namespaces:
+
  * 2X:XX:00:a0:98:XX:XX:XX
  * 2X:XX:00:a0:b8:XX:XX:XX
  * 2X:XX:d0:39:ea:XX:XX:XX
```

```
+
The user must ensure that the target name is not in use elsewhere outside
the cluster. ONTAP cannot verify that the target name is unique outside
the cluster if ONTAP did not generate the target name.

``[-status-admin {down|up}]`` - Administrative Status::
Specifies the administrative status of the FCP service of a Vserver. If
you set this parameter to ``_up_`` , the FCP service will accept login
requests from FCP initiators. If you set this parameter to ``_down_`` ,
FCP initiators will not be allowed to log in.

``[-f, -force <true>]`` - Force [.advaced-priv]#(privilege: advanced)#::
Allows you to specify a World Wide Node Name outside one of the known
vendor registered namespaces.If you use this parameter without a value, it
is set to ``_true_`` , and the command does not error when the specified
WWNN is outside one of the vendor registered namespaces.

== Examples
```

cluster1::> vserver fcp create -vserver vs_1

```
Creates an FCP service on Vserver ``_vs_1_`` .




[[ID1e1e286b4e1b89771596c21439dc4f5e]]
= vserver fcp delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete FCP service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Deletes an FCP service of a Vserver. Before you can delete an FCP service,
the administration status must be ``_down_`` . Use the
xref:{relative_path}vserver-fcp-modify.html[vserver fcp modify] command to
change the administration status.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the FCP service.

== Examples
```

cluster1::> vserver fcp delete -vserver vs_1

```
Deletes the FCP service on Vserver ``_vs_1_`` .

== Related Links

* xref:{relative_path}vserver-fcp-modify.html[vserver fcp modify]




[[IDccc6baf7d0f8746c7eaf766a9c4fff1d]]
= vserver fcp modify
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify FCP service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command modifies an FCP service configuration on a Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the FCP service.

``[-target-name <text>]`` - Target Name [.advaced-priv]#(privilege:
advanced)#::
The FCP World Wide Node Name (WWNN) for the service. All FCP LIFs in the
Vserver will share the specified WWNN. The format for a WWNN is
"XX:XX:XX:XX:XX:XX:XX:XX" where X is a hexadecimal digit.
+
Unless the `force` option is also provided, the specified WWNN must be
within one of the following vendor registered namespaces:
+
  * 2X:XX:00:a0:98:XX:XX:XX
  * 2X:XX:00:a0:b8:XX:XX:XX
  * 2X:XX:d0:39:ea:XX:XX:XX



+
The user must ensure that the target name is not in use elsewhere outside
the cluster. ONTAP cannot verify that the target name is unique outside
the cluster if ONTAP did not generate the target name.

``[-status-admin {down|up}]`` - Administrative Status::
Specifies the administrative status of the FCP service of a Vserver. If
you set this parameter to ``_up_`` , the FCP service accepts login
requests from FCP initiators. If you set this parameter to ``_down_`` ,
FCP initiators cannot log in.

``[-f, -force <true>]`` - Force [.advaced-priv]#(privilege: advanced)#::
Allows you to specify a World Wide Node Name outside one of the known
```

```
vendor registered namespaces.If you use this parameter without a value, it
is set to ``_true_`` , and the command does not error when the specified
WWNN is outside one of the vendor registered namespaces.

== Examples
```

cluster1::> vserver fcp modify -vserver vs_1 -status-admin down

```
Changes the administration status of the FCP service on Vserver ``_vs_1_``
to ``_down_`` .




[[IDe9a0635900a5fe1cdeea04c508aaa99d]]
= vserver fcp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display FCP service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Displays the current status of the FCP service in a cluster.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the FCP services that match the Vserver that
you specify.
```

```
``[-target-name <text>]`` - Target Name::
Use this parameter to display the FCP service that matches the target name
that you specify.

``[-status-admin {down|up}]`` - Administrative Status::
Use this parameter to display the FCP services that match the
administrative status that you specify.

== Examples
```

cluster1::> vserver fcp show
Status
server Target Name Admin
--------- ---------- ------
s0 20:00:00:a0:98:0c:b0:eb up
s2 20:01:00:a0:98:0c:b0:eb up
entries were displayed.

```
Displays the FCP configuration for all the Vservers in the cluster.




[[IDeccd140c04836ac020f86512d5009652]]
= vserver fcp start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Starts the FCP service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command starts the FCP service of a Vserver. When you start the FCP
service, the logical interfaces are brought online.

You must have a license before you can start the FCP service. Use
xref:{relative_path}system-license-add.html[system license add] to enable
the FCP license.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the FCP service.

== Examples
```

cluster1::> vserver fcp start -vserver vs_1
(vserver fcp start)

```
Starts FCP service for Vserver ``_vs_1_`` .

== Related Links

* xref:{relative_path}system-license-add.html[system license add]




[[ID8ac1e19cfaa9c1842bbedbf74f8c556a]]
= vserver fcp stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stops the FCP service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command stops the FCP service of a Vserver. When you stop the FCP
service, the operation status of all FCP logical interfaces in the Vserver
will be ``_down_`` .

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the FCP service.

== Examples
```

cluster1::> vserver fcp stop -vserver vs_1
(vserver fcp stop)

```
Stops FCP service on Vserver ``_vs_1_`` .




[[ID0bff93df5c7074a01ab2d14c12685643]]
= vserver fcp initiator show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display FCP initiators currently connected

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays information about FCP initiators that are currently
logged in.

If you do not specify a Vserver, the command displays all initiators
logged into all FCP Vservers within a cluster. If you specify a Vserver
but not a logical interface, the command displays information about all
initiators connected to all logical interfaces within the specified
Vserver.

If an initiator belongs to an initiator group or has a World Wide Port
Name (WWPN) alias, the command displays this information.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the FCP initiators logged into the Vserver
that you specify.

``[-lif <lif-name>]`` - Logical Interface::
Use this parameter to display the FCP initiators that match the logical
interfaces that you specify.

``[-wwpn <FC WWN>]`` - WWPN::
Use this parameter to display the FCP initiators that matches the World
Wide Port Name (WWPN) that you specify.
```

```
``[-wwnn <FC WWN>]`` - WWNN::
Use this parameter to display the FCP initiator that matches the World
Wide Node Name (WWNN) that you specify.

``[-port-address <Hex Integer>]`` - Port Addres::
Use this parameter to display FCP initiators that match the port address
that you specify.

``[-alias <text>,...]`` - Initiator WWPN Alias::
Use this parameter to display the FCP initiator that matches the alias
name that you specify.

``[-igroup <text>,...]`` - Igroup Name::
Use this parameter to display the FCP initiator that matches the initiator
group that you specify.

``[-data-protocol {fcp|fc-nvme}]`` - Data Protocol::
Use this parameter to display the FCP initiator that matches the data
protocol that you specify.

== Examples
```

cluster1::> vserver fcp initiator show
Logical Initiator Initiator
Vserver Interface WWNN WWPN Igroup
--------- ---------- ---------- ---------- ----------
vs1 vs1.fcp 2f:a2:00:a0:98:0b:56:13
2f:a2:00:a0:98:0b:56:15
igroup1

```
Displays information regarding all logged in FCP initiators.




[[ID8cbf8c9ef8acb9fbb768c32aad755051]]
= vserver fcp interface show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display configuration information for an FCP interface
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays FCP logical interface information. If you do not
specify a Vserver, the command displays all of the FCP data interfaces of
a cluster.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter with other options to display information about FCP
logical interfaces scoped to the specified Vserver.

``[-lif <lif-name>]`` - Logical Interface::
Use this parameter to display FCP logical interfaces that match the names
of logical interfaces that you specify. You can provide a partial logical
interface name, and press tab to complete the name or the closest match.

``[-wwpn <text>]`` - WWPN::
Use this parameter to display FCP logical interfaces that match the World
Wide Port Name (WWPN) that you specify.

``[-wwnn <text>]`` - WWNN::
Use this parameter to display FCP logical interfaces that match the World
Wide Node Name (WWNN) that you specify.

``[-status-admin {up|down}]`` - Administrative Status::
Specifies the configured status of the FCP logical interface. If you set
this parameter to ``_up_`` the command displays all FCP logical interfaces
with the administrative status of ``_up_`` If you set this parameter to
``_down_`` the command displays all the FCP logical interfaces with the
administrative status of ``_down_`` .

``[-status-oper {up|down}]`` - Operational Status::
Specifies the current status of the FCP logical interface. If you set this
parameter to ``_up_`` the command displays all the FCP logical interfaces
```

```
with the operational status of ``_up_`` If you set this parameter to
``_down_`` the command displays all the FCP logical interfaces with the
operational status of ``_down_`` .

``[-status-extended <text>]`` - Extended Status::
Use this parameter to display more detailed information on the status of
the FCP logical interface that you specify.

``[-port-address <Hex Integer>]`` - Host Port Address::
Use this parameter to display FCP logical interfaces that match the host
port address that you specify.

``[-curr-node <nodename>]`` - Current Node::
Use this parameter to display FCP logical interfaces that are on the node
that you specify.

``[-curr-port {<netport>|<ifgrp>}]`` - Current Port::
Use this parameter to display FCP logical interfaces that are on the port
that you specify.

``[-is-home {true|false}]`` - Is Home::
Specifies whether the node hosting the FCP interface is the initially
configured node. If you use this command without using this parameter, it
is set to true, and the command displays all FCP interfaces that are on
the initially configured node.

``[-relative-port-id <integer>]`` - Relative Port ID::
Use this parameter to display FCP logical interfaces that match the
relative target port ID that you specify. The system assigns each LIF and
target portal group a relative target port ID that is Vserver unique. You
cannot change this ID.

== Examples
```

cluster1::> vserver fcp interface show
Logical Status Current Current Is
Vserver Interface Admin/Oper WWPN Node Port Home
--- --- --- --- --- --- ---
vs1 vs1.fcp up/down 2f:a2:00:a0:98:0b:56:13
node1 0c true

```
Displays all FCP interface information.
```

```
[[IDb0378ed4d28c33e2d0ca72bc09c7cd13]]
= vserver fcp nameserver show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display FCP fabric name server entries

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command show entries in the fabric name server database.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to select the Vservers that contain FCP LIFs.

``[-lif <text>]`` - LIF Name::
Use this parameter to select the FCP LIFs.

``[-port-id <Hex Integer>]`` - Port Identifier::
Use this parameter to select the assigned port identifier of the LIF.

``[-unzoned <true>]`` - Show unzoned name server entries::
Use this parameter to show unzoned name server information.

``[-port-type <text>]`` - Port Type::
Use this parameter to select the port type of the LIF.

``[-port-wwn <text>]`` - Port WWN::
Use this parameter to select World Wide Port Name (WWPN) of the LIF.

``[-fabric-port-wwn <text>]`` - Fabric Port WWN::
```

```
Use this parameter to select the fabric World Wide Port Name (WWPN) of the
lif.

``[-node-wwn <text>]`` - Node WWN::
Use this parameter to select the World Wide Node Name (WWNN) of the LIF.

``[-service-class <text>]`` - Service Class::
Use this parameter to select the registered class of services as defined
in the FC-FS standard.

``[-fc4-type <text>]`` - FC4 Type::
Use this parameter to select the registered FC4 type.

``[-switch-port <text>]`` - Switch Port::
Use this parameter to select the name of switch port connected to target
array.

== Examples
```

cluster1::> vserver fcp nameserver show
Port Port FC4
Vserver:Lif Node WWN, Port WWN Id Type Type
--- --- --- --- ---
vs1 :lif1
20:00:00:a0:98:55:73:38
20:01:00:a0:98:55:73:38 8130561 N-Port FCP
20:00:00:90:fa:73:12:dd
10:00:00:90:fa:73:12:dd 8194560 N-Port
vs1 :lif2
20:00:00:90:fa:94:29:ee
10:00:00:90:fa:94:29:ee 8201984 N-Port FCP
3 entries were displayed.

```
[[ID46921f6cffdd6238c6ca96769070feaa]]
= vserver fcp ping-igroup show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Ping FCP by Igroup

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
This command performs a connectivity check (ping) between the FCP
initiators of an igroup and the FCP LIFs for which they are configured.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to select the Vservers that contain initiators and FCP
LIFs.

``[-igroup <text>]`` - Igroup Name::
Use this parameter to select the FCP initiators that belong to the
specified igroup and FCP LIFs that belong to the portset that is bound to
the igroup. If the igroup is not bound to a portset, then the default
portset (all FCP LIFs in the Vserver), is used.

``[-wwpn <text>]`` - FCP initiator WWPN::
Use this parameter to select the FCP initiator WWPN.

``[-lif <text>]`` - LIF Name::
Use this parameter to limit the test to a subset of the FCP LIFs available
for the igroup.

``[-portset <text>]`` - Portset::
Use this parameter to select igroups bound to the specified portset.

``[-node <nodename>]`` - Node::
Use this parameter to select the nodes that contain the specified FCP
LIFs.

``[-status {unknown|reachable|not-reachable|not-zoned|cannot-ping-same-
wwpn|fcp-service-busy|lif-is-down|zone-info-not-available}]`` - Ping
Status::
Use this parameter to select the status of FCP ping command.

``[-ext-status {logged-in|not-logged_in|not-in-fabric|not-in-same-
zone|fabric-info-not-available}]`` - Extended Status::
Use this parameter to select the extended status of FCP ping command.
```

```
``[-check-fabric <true>]`` - Query Fabric Records [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to query the unzoned name server for the FCP initiator
WWPN.

== Examples
```

cluster1::> vserver fcp ping-igroup show
Igroup Node Logical Ping Extended
Vserver Name WWPN Name Interface Status Status
--------- ---------- ---------- --------- --------- --------- ----------
vserver_1
igroup_1 c0:03:ff:e4:70:06:00:e4
node_1
lif_1 reachable wwpn-logged-in
igroup_1 c0:03:ff:e4:70:06:00:e6
node_2
lif_2 not-zoned -
2 entries were displayed.

```
[[IDfdb9b1b7fd340eb40e04f194b40e5605]]
= vserver fcp ping-initiator show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Ping FCP initiator

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command performs a connectivity check (ping) between FCP initiators
and FCP LIFs.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to select the Vservers that contain FCP initiators and
FCP LIFs.

``-wwpn <text>`` - Remote WWPN::
Use this parameter to select the remote WWPN (most likely, FCP initiator).

``[-lif <text>]`` - LIF Name::
Use this parameter to limit the test to a subset of the FCP LIFs available
for the igroup.

``[-check-fabric <true>]`` - Query Fabric Records [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to query the unzoned name server for the FCP initiator
WWPN.

``[-node <nodename>]`` - Node::
Use this parameter to select the nodes that contain the specified FCP
LIFs.

``[-status {unknown|reachable|not-reachable|not-zoned|cannot-ping-same-
wwpn|fcp-service-busy|lif-is-down|zone-info-not-available}]`` - Ping
Status::
Use this parameter to select the result of FCP ping command.

``[-ext-status {logged-in|not-logged_in|not-in-fabric|not-in-same-
zone|fabric-info-not-available}]`` - Extended Status::
Use this parameter to select the extended result of FCP ping command.

== Examples
```

cluster1::> vserver fcp ping-initiator show
Node Logical Ping Extended
Vserver WWPN Name Interface Status Status
--------- ---------- --------- --------- --------- ----------
vserver_1
c0:03:ff:e4:70:06:00:e4
node_1
lif_1 reachable wwpn-logged-in
c0:03:ff:e4:70:06:00:e6
node_2
lif_2 not-zoned -
2 entries were displayed.

```
[[ID07d8077124d5b6a392d54c222922ce32]]
= vserver fcp portname set
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Assigns a new WWPN to a FCP adapter

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command assigns a new World Wide Port Name (WWPN) to a logical
interface. The administration status of logical interface must be down
before you can change the WWPN.

Use the xref:{relative_path}network-interface-modify.html[network
interface modify] to change the administration status of the logical
interface.

== Parameters

``-vserver <Vserver Name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Specifies the Vserver.

``-lif <lif-name>`` - Logical Interface [.advaced-priv]#(privilege:
advanced)#::
Specifies the logical interface to which you want to assign a new WWPN.

``-wwpn <text>`` - FCP Adapter WWPN [.advaced-priv]#(privilege:
advanced)#::
Specifies the WWPN that you want to change.

``[-f, -force <true>]`` - Force [.advaced-priv]#(privilege: advanced)#::
Allows you to use a WWPN that is not in the format 2X:XX:0a:09:80:XX:XX:XX
when set to true. If you use this parameter without a value, it is set to
true, and the command does not prompt you when the WWNN does not follow
this format.

== Examples
```

cluster1::*> vserver fcp portname set -vserver vs_1 -lif vs_1.fcp -wwpn 2f:a2:00:a0:98:0b:56:13

```
Sets a new WWPN for LIF vs_1.fcp on Vserver vs_1.

== Related Links

* xref:{relative_path}network-interface-modify.html[network interface
modify]




[[IDaf1f7a416af64e7e95f18974e63a77d2]]
= vserver fcp portname show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display WWPN for FCP logical interfaces

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays a list of World Wide Port Names (WWPN) that are used
by the FCP logical interfaces.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display a list of FCP logical interfaces and their
WWPNs that match the Vserver name you specify.

``[-lif <lif-name>]`` - Logical Interface::
```

```
Use this parameter to display a list of FCP logical interfaces and their
WWPNs that match the logical interface that you specify. You can use
wildcards in the logical interface to display a specific group of logical
interfaces.

``[-wwpn <text>]`` - WWPN::
Use this parameter to display a list of FCP logical interfaces and their
WWPNs that match the WWPN that you specify. You can use wildcards in the
WWPN to display a specific group of WWPNs.

== Examples
```

cluster1::> vserver fcp portname show
Logical
Vserver Interface WWPN
--------- -------------- ------------------------
vs_a vs_a.fcp 2f:a2:00:a0:98:0b:56:13
vs_io1 vs_io1.fcp 2f:9e:00:a0:98:0b:56:13
vs_2 lif2 2f:a3:00:a0:98:0b:56:13
vs_2 lif3 2f:a4:00:a0:98:0b:56:13
vs_2 lif4 2f:a5:00:a0:98:0b:56:13
vs_2 lif5 2f:a6:00:a0:98:0b:56:13
vs_2 vs_2.fcp 2f:9a:00:a0:98:0b:56:13
vs1 vs1.fcp 2f:9d:00:a0:98:0b:56:13
vs1 vs1.fcp2 2f:97:00:a0:98:0b:56:13

```
Displays the WWPNs for each FCP logical interface for all the Vservers in
a cluster.




[[IDc77bda2f5f0baf549766b9992f9964eb]]
= vserver fcp topology show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show FCP topology interconnect elements

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
Show FCP topology interconnect elements

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to select the interconnect elements for the specified
Vservers.

``[-lif <text>]`` - LIF Name::
Use this parameter to select the interconnect elements for the specified
FCP LIFs.

``[-domain-id <integer>]`` - Domain Identifier::
Use this parameter to select the interconnect elements with the specified
domain identifier

``[-logical-name <text>]`` - Logical Name::
Use this parameter to select the interconnect elements with the specified
logical name

``[-vendor <text>]`` - Vendor::
Use this parameter to select the interconnect elements with the specified
vendor

``[-release <text>]`` - Release::
Use this parameter to select the interconnect elements with the specified
release

``[-wwn <text>]`` - World Wide Name::
Use this parameter to select the interconnect elements with the specified
World Wide Name

``[-port-count <integer>]`` - Port Count::
Use this parameter to select the interconnect elements with the specified
port count

== Examples
```

```
cluster1::> vserver fcp topology show
Domain Logical Port
Vserver Lif Name Id Name WWN Count
------- -------------- ------ ------------------ ---------------------- -----
vs1 lif1
98 ssan-fc0e-fit-01 20:05:00:05:73:bd:a3:01 15
99 ssan-fc0e-d38 20:05:8c:60:4f:04:f1:01 11
112 ssan-fc0e-5 20:05:00:0d:ec:ca:0b:41 19
119 ssan-fc0e-core-a 20:05:54:7f:ee:02:c1:01 18
159 ssan-fc0e-7 20:05:00:05:9b:24:6e:c1 38
169 sdev-fc0e-gg26 20:05:54:7f:ee:31:06:81 53
174 ssan-fc0e-d46 20:05:00:05:9b:7d:f8:01 16
177 ssan-fc0e-e49 20:05:54:7f:ee:ef:1c:81 19
180 ssan-fc0e-d40 20:05:00:05:9b:79:a3:c1 20
vs1 lif2
229 ssan-fc0e-6 20:05:00:05:73:c8:8f:01 33
233 ssan-fc0e-e45 20:05:54:7f:ee:a0:67:01 8
11 entries were displayed.
```

```
The example above show FCP topology interconnect information for the
cluster.




[[ID7ac55e7828662cca773852949d715fef]]
= vserver fcp wwn blacklist show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays the blacklisted WWNs

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command displays WWNs that have been blacklisted from re-use.

A blacklisted WWN is a WWN that is prohibited for use as either a fiber
channel protocol service WWNN or a fiber channel data LIF WWPN.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-wwn <text>]`` - World Wide Name [.advaced-priv]#(privilege:
advanced)#::
Selects the blacklisted WWNs that match the parameter value.

``[-vserver <Vserver Name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Selects the blacklisted WWNs that were previously assigned to the
Vserver(s) that match the parameter value.

== Examples
```

cluster1::> vserver fcp wwn blacklist show
WWN Vserver
---------- ----------
01:02:03:04:05:06:07:08 vs1
01:02:03:04:05:06:07:09 vs1
2 entries were displayed.

```
Displays the blacklisted WWNs.




[[IDc1527b2d50eb730410fdb340e194bae2]]
= vserver fcp wwpn-alias remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Removes an alias for a World Wide Port Name of an initiator.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command removes an alias from a World Wide Port Name (WWPN).

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

{ ``-a, -alias <text>,...`` - Initiator WWPN Alias::
Specifies the alias of the WWPN that you want to remove.

| ``-w, -wwpn <FC WWN>`` - Initiator WWPN }::
Specifies the WWPN.

== Examples
```

cluster1::> vserver fcp wwpn-alias remove -vserver vs_1 -wwpn 2f:a0:00:a0:98:0b:56:13

```
On Vserver ``_vs_1_`` , removes all the aliases on WWPN
``_2f:a0:00:a0:98:0b:56:13_`` .
```

cluster1::> vserver fcp wwpn-alias remove -vserver vs_1 -alias my_alias

```
On Vserver ``_vs_1_`` , removes the alias ``_my_alias_`` .




[[ID8b0daf19e7654be8a99be6fab3224085]]
= vserver fcp wwpn-alias set
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Set an alias for a World Wide Port Name of an initiator that might login
to the target.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates a new alias for a World Wide Port Name (WWPN). You
can create multiple aliases for a WWPN, but you cannot use the same alias
for multiple WWPNs.

An alias name is a case-sensitive name that must contain one to 32
characters. Spaces are not allowed.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-a, -alias <text>`` - Initiator WWPN Alias::
Specifies the alias of the WWPN.

``-w, -wwpn <FC WWN>`` - Initiator WWPN::
Specifies the WWPN.

``[-f, -force <true>]`` - Force::
Allows you to override a WWPN associated with an existing alias with a
newly specified WWPN. If you use this parameter without a value, it is set
to true, and the command does not prompt you when you override an existing
alias.

== Examples
```

cluster1::> vserver fcp wwpn-alias set -vserver vs_1 -alias my_alias -wwpn 2f:a0:00:a0:98:0b:56:13

```
Sets the alias ``_my_alias_`` for the WWPN ``_2f:a0:00:a0:98:0b:56:13_`` .




[[ID42e0c37eb749d8ed26a39c66166e0938]]
= vserver fcp wwpn-alias show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Displays a list of the WWPN aliases configured for initiators

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays aliases associated with World Wide Port Names
(WWPN).

NOTE: You can also use these commands to display WWPN aliases:

* xref:{relative_path}lun-igroup-show.html[lun igroup show]
* xref:{relative_path}lun-igroup-create.html[lun igroup create]
* xref:{relative_path}lun-igroup-add.html[lun igroup add]
* xref:{relative_path}lun-igroup-remove.html[lun igroup remove]
* xref:{relative_path}vserver-fcp-show.html[vserver fcp show]




== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display a list of WWPNs and the associated aliases
that match the Vserver name that you specify.

``[-a, -alias <text>]`` - Initiator WWPN Alias::
Use this parameter to display the WWPN that matches the alias that you
specify.

``[-w, -wwpn <FC WWN>]`` - Initiator WWPN::
Use this parameter to display a list of aliases that match the WWPN that
you specify.

== Examples
```

cluster1::> vserver fcp wwpn-alias show
Initiator Initiator
Vserver WWPN Alias
--------- ---------- ----------
vs1 2f:a0:00:a0:98:0b:56:13 my_alias

```
Displays the alias my_alias for the WWPN 2f:a0:00:a0:98:0b:56:13 on
Vserver vs1.

== Related Links

* xref:{relative_path}lun-igroup-show.html[lun igroup show]
* xref:{relative_path}lun-igroup-create.html[lun igroup create]
* xref:{relative_path}lun-igroup-add.html[lun igroup add]
* xref:{relative_path}lun-igroup-remove.html[lun igroup remove]
* xref:{relative_path}vserver-fcp-show.html[vserver fcp show]




:leveloffset: -1


= vserver fpolicy commands

:leveloffset: +1


[[ID476e8fc30d33ccd2b316abad571dfe75]]
= vserver fpolicy disable
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy disable` command disables an FPolicy policy for the
specified Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
disable an FPolicy policy.

``-policy-name <Policy name>`` - Policy::
 This parameter specifies the name of the FPolicy policy you want to
disable.

== Examples

The following command disables an FPolicy policy.
```

cluster1::> vserver fpolicy show
Vserver Policy Name Sequence Status Engine
--- --- --- --- ---
vs1.example.com vs1_pol - off native
vs2.example.com vs2_pol 5 on external
2 entries were displayed.

cluster1::> vserver fpolicy disable -vserver vs2.example.com -policy-name vs2_pol

cluster1::> vserver fpolicy show
Vserver Policy Name Sequence Status Engine
--- --- --- --- ---
vs1.example.com vs1_pol - off native
vs2.example.com vs2_pol - off external
2 entries were displayed.

```
[[ID37df9e28c8de40ec57c5e8e7b8b39de2]]
= vserver fpolicy enable
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy enable` command enables FPolicy policies for the
specified Vserver and sets their sequence (priority). The sequence is used
when multiple policies have subscribed to the same file access event. To
modify the sequence number of a policy, the administrator must disable the
policy (if it is enabled) and then use this command to enable it with the
new sequence number. Policies that use the ``_native_`` engine
configuration have a higher priority than policies for any other engine,
regardless of the sequence number assigned to them.

NOTE: Events on FlexGroup volumes do not notify the FPolicy server.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
enable an FPolicy policy. The Vserver administrator can enable FPolicy
policies created within the scope of the Vserver and can also enable an
FPolicy policy created by the cluster administrator. The cluster
administrator can enable FPolicy policies for any Vserver but cannot
enable them with a scope of cluster. The scope is determined at a Vserver
level.

``-policy-name <Policy name>`` - Policy::
 This parameter specifies the name of the FPolicy policy you want to
enable.

``-sequence-number <integer>`` - Policy Sequence Number::
This parameter specifies the sequence number that is assigned to the
policy.

== Examples
```

```
The following command enables an FPolicy policy:
```

cluster1::> vserver fpolicy show
Vserver Policy Name Sequence Status Engine
---------- ---------- -------- ------ --------
vs1.example.com vs1_pol - off native
vs2.example.com vs2_pol - off external
2 entries were displayed.
cluster1::> vserver fpolicy enable -vserver vs2.example.com -policy-name vs2_pol -sequence-number 5

cluster1::> vserver fpolicy show
Vserver Policy Name Sequence Status Engine
---------- ---------- -------- ------- -------
vs1.example.com vs1_pol - off native
vs2.example.com vs2_pol 5 on external
2 entries were displayed.

```
[[ID4787ec7bf681d0d059bebc41edbad44d]]
= vserver fpolicy engine-connect
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Establish a connection to FPolicy server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy engine-connect` command connects an FPolicy server to
a specified node. Connecting the FPolicy server to a node enables FPolicy
processing, providing the FPolicy configuration is complete. Before
connecting an FPolicy server to a node, you must configure FPolicy by
completing the following tasks:

*  Create an FPolicy event
*  Create an FPolicy external engine
*  Create an FPolicy policy
*  Create a scope for the FPolicy policy



NOTE: The FPolicy event and external engine must be attached to the
FPolicy policy.
```

```
NOTE: The FPolicy policy should be enabled.



== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node that you want to connect to the FPolicy
server. The value local specifies the current node.

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver that you want to connect to the
specified FPolicy server using the specified FPolicy policy.

``-policy-name <Policy name>`` - Policy::
 This parameter specifies the name of the FPolicy policy that is attached
to an external engine.

``-server <IP Address>`` - Server::
This parameter specifies the FPolicy server to which you want to connect
the node. The specified server must be present in the external engine
configuration of the above specified policy.

== Examples

The following example connects an FPolicy server.
```

cluster1::> vserver fpolicy engine-connect -node FPolicy-01 -vserver vs1.example.com -policy-name p -server 1.1.1.1
cluster1::> vserver fpolicy show
FPolicy Server- Server-
Vserver Policy Node Server status type
--- --- --- --- --- ---
vs1.example.com p FPolicy-01 1.1.1.1 connected primary

```
[[ID1ee3a4c57fbcecd4d3cf670b2b1898f9]]
= vserver fpolicy engine-disconnect
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Terminate connection to FPolicy server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy engine-disconnect` command disconnects an FPolicy
server from a specified node.

== Parameters

``-node {<nodename>|local}`` - Node::
This parameter specifies the node that you want to disconnect from the
FPolicy server. The value local specifies the current node.

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver that you want to disconnect from the
specified FPolicy server with the specified attached FPolicy policy.

``-policy-name <Policy name>`` - Policy::
 This parameter specifies the name of the FPolicy policy that is attached
with an external engine.

``-server <IP Address>`` - Server::
This parameter specifies the FPolicy server you want to disconnect. The
specified server must be present in the external engine configuration of
the above specified FPolicy policy.

== Examples

The following example disconnects an FPolicy server.
```

cluster1::> vserver fpolicy engine-disconnect -node FPolicy-01 -vserver vs1.example.com -policy-name p
-server 1.1.1.1
cluster1::> vserver fpolicy show
FPolicy Server- Server-
Vserver Policy Node Server status type

--- --- --- --- --- ---
vs1.example.com p FPolicy-01 1.1.1.1 disconnected primary

```
[[ID04514d67d5409ac087d6304ac8cc349d]]
= vserver fpolicy show-enabled
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display all enabled policies

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver fpolicy show-enabled` command displays information about all
enabled policies in the Vserver. The command output depends on the
parameter or parameters specified with the command. If you do not specify
any parameters, the command displays the following information about all
FPolicy policies:

* Vserver name
* Policy name
* Priority



You can specify the `-fields` parameter to specify which fields of
information to display about FPolicy policies.

You can specify the `-instance` parameter to display information for all
FPolicy policies in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.
```

```
``[-vserver <vserver>]`` - Vserver::
If you specify this parameter, the command displays information only about
the FPolicy policies for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy Name::
If you specify this parameter, the command displays information only about
the FPolicy policy that you specify.

``[-priority <text>]`` - Policy Priority::
If you specify this parameter, the command displays information only about
the FPolicy policies with the policy priority that you specify.

== Examples

The following example displays the information about enabled FPolicy
policies on the cluster.
```

cluster1::> vserver fpolicy show-enabled
Vserver Policy Name Priority

--- --- ---
vs1.example.com pol_native native
vs1.example.com pol_native2 native
vs1.example.com pol1 2
vs1.example.com pol2 4

```
[[ID26169c56670b9da29acadf5bb6727675]]
= vserver fpolicy show-engine
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display FPolicy server status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy show-engine` command displays status information for
all FPolicy external engines or displays status information only for
FPolicy servers for a specified Vserver. The command output depends on the
parameter or parameters specified with the command. If you do not specify
any parameters, the command displays the following information for all
```

```
FPolicy servers:

* Vserver name
* Node name
* FPolicy policy name
* FPolicy server IP Address
* FPolicy server status
* FPolicy server type



You can specify the `-fields` parameter to specify which fields of
information to display about FPolicy servers. You can specify specific
parameters to display only information that matches those parameters. For
instance, to display information only about all FPolicy servers (external
engines) that are connected, run the command with the `-fields` parameter
set to server and `-server-status ` parameter set to connected.

You can specify the `-instance` parameter to display all information for
all policies in the list form.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-node {<nodename>|local}]`` - Node::
 If you specify this parameter, the command displays information only
about the FPolicy external engine attached to the specified node.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the FPolicy server for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
the FPolicy servers that are attached with the specified policy.

``[-server <IP Address>]`` - Server::
If you specify this parameter, the command displays information only about
the FPolicy servers that you specify.
```

```
``[-server-status <Status>]`` - Server Status::
If you specify this parameter, the command displays information only about
the FPolicy servers that have the specified status.

``[-server-type <Server Type>]`` - Server Type::
If you specify this parameter, the command displays information only about
the FPolicy servers that have the specified server type.

``[-connected-since <MM/DD/YYYY HH:MM:SS>]`` - Time FPolicy Server was
Connected::
If you specify this parameter, the command displays information only about
the FPolicy servers that have been connected since the specified time.

``[-disconnected-since <MM/DD/YYYY HH:MM:SS>]`` - Time FPolicy Server was
Disconnected::
If you specify this parameter, the command displays information only about
the FPolicy servers that have been disconnected since the specified time.

``[-disconnect-reason <text>]`` - Reason for FPolicy Server
Disconnection::
If you specify this parameter, the command displays information only about
the FPolicy servers that are disconnected because of the specified reason.

``[-disconnect-reason-id <integer>]`` - ID for FPolicy Server
Disconnection::
If you specify this parameter, the command displays information about the
FPolicy servers that are disconnected because of the specified disconnect
reason ID. There is a unique ID associated with each disconnect reason,
which can be used to identify the reason for FPolicy server disconnection.

``[-session-id <text>]`` - Session ID::
If you specify this parameter, the command displays information about the
FPolicy server that is connected with the specified session ID. There is a
unique session ID associated with each connection to FPolicy server, which
can be used to identify the established connection.

== Examples

This example displays information about all FPolicy servers (external
engines).
```

cluster1::> vserver fpolicy show-engine
FPolicy Server- Server-
Vserver Policy Node Server status type
--- --- --- --- --- ---
vs2.example.com vs2_pol FPolicy-01 9.9.9.9 connected primary
vs1.example.com vs1_pol FPolicy-01 1.1.1.1 connected primary

2 entries were displayed.

```
This example displays information only about all connected FPolicy servers
(external engines).
```

cluster1::> vserver fpolicy show-engine -fields server -server-status connected
node vserver policy-name server
--- --- --- ---
FPolicy-01 vs1.example.com vs1_pol 1.1.1.1

```
This example displays information about an FPolicy server.
```

cluster1::> vserver fpolicy show-engine -server 10.72.204.118 -instance
Node: fpol-01
Vserver: vserver_1.example.com
Policy: pol_cifs
Server: 10.72.204.118
Server Status: disconnected
Server Type: primary
Time FPolicy Server was Connected: -
Time FPolicy Server was Disconnected: 2/5/2013 05:06:22
Reason for FPolicy Server Disconnection: TCP Connection to FPolicy server failed.
ID for FPolicy Server Disconnection: 9307
Session ID:

```
[[ID9bd93586454ca60172803ae57eb53751]]
= vserver fpolicy show-passthrough-read-connection
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display connection status for FPolicy passthrough-read

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy show-passthrough-read-connection` command displays
the status of the passthrough-read connection from all FPolicy servers.
Passthrough-read is a way to read data for offline files without restoring
the files to primary storage. If you do not specify any parameters, the
command displays following information about the passthrough-read
connection from FPolicy servers:
```

```
* Vserver name
* FPolicy policy name
* Node name
* FPolicy server IP address
* Passthrough-read connection status



You can specify the `-fields` parameter to specify which fields of
information to display. In addition to the fields above, you can display
the following fields.

* Session ID of the control channel
* Time passthrough-read channel was connected
* Time passthrough-read channel was disconnected
* Reason for passthrough-read channel disconnection



You can specify the `-instance` parameter to display information for all
passthrough-read connections in the list form.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-node {<nodename>|local}]`` - Node::
 If you specify this parameter, the command displays information only
about the passthrough-read connections on the specified node.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the passthrough-read connections for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
the passthrough-read connections that are attached with the specified
FPolicy policy.

``[-server <IP Address>]`` - Server::
```

```
If you specify this parameter, the command displays information only about
the passthrough-read connections from the specified FPolicy server.

``[-control-session-id <text>]`` - Session ID of the Control Channel::
If you specify this parameter, the command displays information only about
the passthrough-read connections that are connected with the specified
control session ID. The passthrough-read connection is attached to a
control connection that has a unique control session ID.

``[-server-status <Status of fpolicy passthrough-read connection>]`` -
Server Status::
If you specify this parameter, the command displays information only about
the passthrough-read connections that have the specified status.

``[-connected-since <MM/DD/YYYY HH:MM:SS>]`` - Time Channel Was
Connected::
If you specify this parameter, the command displays information only about
the passthrough-read connections that have the specified connection time.

``[-disconnected-since <MM/DD/YYYY HH:MM:SS>]`` - Time Channel Was
Disconnected::
If you specify this parameter, the command displays information only about
the passthrough-read connections that have the specified disconnection
time.

``[-disconnect-reason <Reason for fpolicy passthrough-read
disconnection>]`` - Reason for Disconnection::
If you specify this parameter, the command displays information only about
the passthrough-read connections that are disconnected because of the
specified disconnect reason.

== Examples

This example displays information about passthrough-read connections from
all FPolicy servers.
```

cluster1::> vserver fpolicy show-passthrough-read-connection
FPolicy Server
Vserver Policy Name Node Server Status
--- --- --- --- ---
vs2.example.com pol_cifs_2 FPolicy-01 2.2.2.2 disconnected
vs1.example.com pol_cifs_1 FPolicy-01 1.1.1.1 connected
2 entries were displayed.

```
This example displays information about passthrough-read connections from
all connected FPolicy servers.
```

cluster1::> vserver fpolicy show-passthrough-read-connection -server-status connected
FPolicy Server
Vserver Policy Name Node Server Status
--- --- --- --- ---
vs1.example.com pol_cifs_1 FPolicy-01 1.1.1.1 connected

```
This example displays information about passthrough-read connections from
FPolicy servers configured in an FPolicy policy.
```

cluster1::> vserver fpolicy show-passthrough-read-connection -policy-name pol_cifs_1 -instance
Node: FPolicy-01
Vserver: vserver_1.example.com
Policy: pol_cifs_1
Server: 2.2.2.2
Session ID of the Control Channel: 8cef052e-2502-11e3-88d4-123478563412
Server Status: connected
Time Passthrough Read Channel was Connected: 9/24/2013 10:17:45
Time Passthrough Read Channel was Disconnected: -
Reason for Passthrough Read Channel Disconnection: none

```
[[IDc903d9aa2853ba0660775fdf3fe25b10]]
= vserver fpolicy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display all policies with status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy show` command displays status information about all
FPolicy policies in the Vserver. The command output depends on the
parameter or parameters specified with the command. If you do not specify
any parameters, the command displays the following information about all
FPolicy policies:

* Vserver name
* Policy name
* Sequence number
* Status
```

```
You can specify the `-fields` parameter to specify which fields of
information to display about FPolicy policies.

You can specify the `-instance` parameter to display information for all
FPolicy policies in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the FPolicy policies for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
the FPolicy policy that you specify.

``[-sequence-number <integer>]`` - Sequence Number::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified sequence-number.

``[-status {on|off}]`` - Status::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified status.

``[-engine <Engine name>]`` - FPolicy Engine::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified engine.

== Examples

The following example displays the information about FPolicy policies on
the cluster using the `vserver fpolicy show` command.
```

cluster1::> vserver fpolicy show
Sequence
Vserver Policy Name Number Status Engine
---------- ---------- -------- -------- ---------
FPolicy cserver_policy - off eng1

vs1.example.com v1p1 - off eng2
vs1.example.com v1p2 - off native
vs1.example.com v1p3 - off native
vs1.example.com cserver_policy - off eng1
vs2.example.com v1p1 3 on native
vs2.example.com v1p2 1 on eng3
vs2.example.com cserver_policy 2 on eng1
8 entries were displayed.

```
[[ID0349283bc140ebc0f447853f1186703e]]
= vserver fpolicy persistent-store create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a Persist Store

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy persistent-store create` command create an FPolicy
Persistent Store entry for the Vserver. This can then be used for enabling
the persistent mode for Fpolicy events by specifying the Fpolicy
Persistent Store name for the parameter "persistent-store" when
creating/modifying the Fpolicy policy. In persistent mode, when the
Persistent Store is full, event notifications are dropped.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
create FPolicy Persistent Store.

``-persistent-store <text>`` - Persistent Store Name::
This parameter specifies the name of the FPolicy Persistent Store that you
want to create.

``-volume <text>`` - Volume name of the Persistent Store::
This parameter specifies volume name for the FPolicy Persistent Store.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of the Persistent Store::
This parameter specifies size of the FPolicy Persistent Store.
```

```
``[-autosize-mode {off|grow|grow_shrink}]`` - Autosize Mode for the
Volume::
This parameter specifies autosize mode of the volume. The valid values
are: ``_off_`` , ``_grow_`` , ``_grow_shrink_`` . The default value is
``_off_`` .

== Examples

The following example creates an FPolicy persistent-store.
```

cluster1::> vserver fpolicy persistent-store create -vserver vs1.example.com -persistent-store ps1 -volume psvol -size 1GB -autosize-mode grow

cluster1::> vserver fpolicy persistent-store show -vserver vs1 -persistent-store ps1
Vserver: vs1.example.com
Persistent Store Name: ps1
Volume name of the Persistent store: psvol
Size of the Persistent Store: 1GB
Autosize Mode for the Volume: grow

```
[[IDc7437066bade22ed952fe06293f469f1]]
= vserver fpolicy persistent-store delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Persist Store

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy persistent-store delete` command deletes an FPolicy
Persistent Store.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver from which you want to
delete the FPolicy Persistent Store.

``-persistent-store <text>`` - Persistent Store Name::
This parameter specifies the name of the FPolicy Persistent Store that you
want to delete.

``[-delete-volume {true|false}]`` - Delete the Persistent Store Volume::
This parameter specifies if the associated volume for the FPolicy
Persistent Store has to be deleted. By default, it will be false.

== Examples

The following example deletes an FPolicy Persistent Store.
```

cluster1::> vserver fpolicy persistent-store delete -vserver vs1.example.com -persistent-store ps1 -delete -volume true

```
[[ID56d46cdfd5615534a8575a4eddd4c387]]
= vserver fpolicy persistent-store modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a Persist Store

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy persistent-store modify` command modifies an FPolicy
Persistent Store.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
modify an FPolicy Persistent Store.

``-persistent-store <text>`` - Persistent Store Name::
 This parameter specifies name of the FPolicy Persistent Store that you
want to modify.

``[-volume <text>]`` - Volume name of the Persistent Store::
This parameter specifies volume name for the FPolicy Persistent Store.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of the Persistent Store::
This parameter specifies size of the FPolicy Persistent Store.

``[-delete-volume {true|false}]`` - Delete the previous Volume::
This parameter specifies if the previously associated volume for the
FPolicy Persistent Store has to be deleted. By default, it will be false.

== Examples

The following example modifies an FPolicy Persistent Store.
```

cluster1::> vserver fpolicy persistent-store modify -vserver vs1.example.com -persistent-store ps1 -volume psvol -size 1GB -delete-volume true

cluster1::> vserver fpolicy persistent-store show -vserver vs1.example.com -persistent-store ps1 -size 1GB

Vserver: vs1.example.com
Persistent Store Name: ps1
Volume name of the Persistent Store: psvol
Size of the Persistent Store: 1GB

```
[[ID1fefc5684ed41d636e5be6480ba88930]]
= vserver fpolicy persistent-store show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Persist Store details

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy persistent-store show` command displays information
about all FPolicy Persistent Store belonging to the Vserver. Any Vserver
administrator can see FPolicy persistent store associated with their
Vserver. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays information about all FPolicy persistent store:

* Vserver Name
* Persistent Store Name
* Persistent Store Volume name



You can specify the `-fields` parameter to specify which fields of
information to display about FPolicy persistent-store. You can specify
additional parameters to display only information that matches those
parameters. For example, to display information only about FPolicy
Persistent Store where the FPolicy persistent store belong to particular
Vserver vs1, run the command with the `-vserver` parameter set to vs1.

You can specify the `-instance` parameter to display all information for
all policies in the list form.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the FPolicy Persistent Store for the specified Vserver.

``[-persistent-store <text>]`` - Persistent Store Name::
If you specify this parameter, the command displays information only about
the FPolicy Persistent Store that you specify.

``[-volume <text>]`` - Volume name of the Persistent Store::
If you specify this parameter, the command displays information only about
the FPolicy Persistent Store that use the specified volume name.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of the Persistent Store::
If you specify this parameter, the command displays information only about
the FPolicy Persistent Store that has the specified size.

``[-autosize-mode {off|grow|grow_shrink}]`` - Autosize Mode for the
Volume::
If you specify this parameter, the command displays information only about
the FPolicy Persistent Store that has the specified autosize mode.

== Examples

The following example displays the information about FPolicy Persistent
Store on the cluster using the `vserver fpolicy persistent-store show`
command.
```

cluster1::> vserver fpolicy persistent-store show

Vserver Persistent Store Volume Size
--- --- --- ---
vs1.example.com ps1 psvol 1GB
vs2.example.com ps2 psvol1 100MB

```
[[ID214d68a0f9b8ddf2032dd2362a184898]]
= vserver fpolicy policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy create` command creates an FPolicy policy. You
must create an FPolicy event name before creating an FPolicy policy. If
you are using an external FPolicy server, you must also create an FPolicy
engine before creating a policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
create an FPolicy policy.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the FPolicy policy that you want to
create. An FPolicy policy name can be up to 256 characters long and is a
string that can only contain any combination of ASCII-range alphanumeric
characters (a-z, A-Z, 0-9), "_" and "." .

``-events <Event name>,...`` - Events to Monitor::
This parameter specifies a list of events to monitor for the FPolicy
policy. All the events in the event list should be created by the
administrator of the specified Vserver or the cluster administrator. The
events must already exist. Create events using the `fpolicy policy event
create` command.

``-engine <Engine name>`` - FPolicy Engine::
This parameter specifies an external engine for this FPolicy policy. An
external engine contains information required by the node to send
notifications to an FPolicy server. The Vserver administrator of the
specified Vserver or the cluster administrator creates the external engine
prior to creating the FPolicy policy. If this parameter is not specified,
the default ``_native_`` external engine is used. The ``_native_``
external engine is internal to ONTAP and is used if you want to configure
native file blocking and you do not want to use an external FPolicy
server.

``[-is-mandatory {true|false}]`` - Is Mandatory Screening Required::
This parameter specifies what action to take on a file access event in a
```

```
case when all primary and secondary servers are down or no response is
received from the FPolicy servers within a given timeout period. When this
parameter is set to ``_true_`` , file access events will be denied under
these circumstances. To allow file access events under these
circumstances, set this parameter to ``_false_`` . By default, it is
``_true_`` .

``[-allow-privileged-access {yes|no}]`` - Allow Privileged Access::
This parameter specifies privileged access for FPolicy servers. It is used
to specify whether privileged access is required for FPolicy servers.
Privileged access is used when the FPolicy server requires direct access
to the cluster nodes. With this option set to ``_yes_`` , FPolicy servers
can access files on the cluster using a separate data channel with
privileged access. By default, it is ``_no_`` .

``[-privileged-user-name <text>]`` - User Name for Privileged Access::
This parameter specifies the privileged user name. It is used to specify
the privileged user name for accessing files on the cluster using a
separate data channel with privileged access. The input for this field
should be in ``_"domain\user name"_`` format. If `-allow-privileged-
access` is set to ``_no_`` , any value set for this field is ignored.

``[-is-passthrough-read-enabled {true|false}]`` - Is Passthrough Read
Enabled::
This parameter specifies whether passthrough-read should be allowed for
FPolicy servers registered for the policy. Passthrough-read is a way to
read data for offline files without restoring the files to primary
storage. Offline files are the files which have been moved to secondary
storage. If passthrough-read is enabled, the FPolicy server provides the
data for the file over a separate channel instead of restoring the file to
primary storage. By default, this parameter is ``_false_`` .

``[-persistent-store <text>]`` - Persistent Store Name::
This parameter specifies persistent storage name. This can then be used
for enabling the Persistent mode for Fpolicy events.

== Examples

The following example creates an FPolicy policy.
```

cluster1::> vserver fpolicy policy create -vserver vs1.example.com -policy-name vs1_pol -events cserver_evt,v1e1
-engine native -is-mandatory true -allow-privileged-access no -is-passthrough-read-enabled false

cluster1::> vserver fpolicy policy show -vserver vs1.example.com -policy-name vs1_pol
Vserver: vs1.example.com
Policy Name: vs1_pol

Events to Monitor: cserver_evt, v1e1
FPolicy Engine: native
Is Mandatory Screening Required: true
Allow Privileged Access: no
User Name for Privileged Access: -
Is Passthrough Read Enabled: false
persistent-store: -

```
[[ID14096489b0f4adf03dd287f89306fa27]]
= vserver fpolicy policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy delete` command deletes an FPolicy policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver from which you want to
delete the FPolicy policy.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the FPolicy policy that you want to
delete.

== Examples

The following example deletes an FPolicy policy.
```

cluster1::> vserver fpolicy policy delete -vserver vs1.example.com -policy-name vs1_pol

```
[[IDae50bb96d62fcdc0e4f5c1b93f40a076]]
= vserver fpolicy policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy modify` command modifies an FPolicy policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
modify an FPolicy policy.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the FPolicy policy that you want to
modify. An FPolicy policy name can be up to 256 characters long and is a
string that can only contain any combination of ASCII-range alphanumeric
characters (a-z, A-Z, 0-9), "_" and ".".

``[-events <Event name>,...]`` - Events to Monitor::
This parameter specifies a list of events to monitor for the FPolicy
policy. All the events in the event list should be created by the
administrator of the specified Vserver or the cluster administrator. The
events must already exist. Create events using the `fpolicy policy event
create` command.

``[-engine <Engine name>]`` - FPolicy Engine::
This parameter specifies an external engine for this FPolicy policy. An
external engine contains information required by the node to send
notifications to an FPolicy server. The Vserver administrator of the
specified Vserver or the cluster administrator creates the external engine
prior to modifying the FPolicy policy. If this parameter is not specified,
the default ``_native_`` external engine is used. The ``_native _``
external engine is internal to ONTAP and is used if you want to configure
``_native_`` file blocking and you do not want to use an external FPolicy
server.

``[-is-mandatory {true|false}]`` - Is Mandatory Screening Required::
This parameter specifies what action to take on a file access event in a
case when all primary and secondary servers are down or no response is
received from the FPolicy servers within a given timeout period. When this
parameter is set to ``_true_`` , file access events will be denied under
```

```
these circumstances. To allow file access events under these
circumstances, set this parameter to ``_false_`` . By default, it is
``_true_`` .

``[-allow-privileged-access {yes|no}]`` - Allow Privileged Access::
This parameter specifies privileged access for FPolicy servers. It is used
to specify whether privileged access is required for FPolicy servers.
Privileged access is used when the FPolicy server requires direct access
to the cluster nodes. With this option set to ``_yes_`` , FPolicy servers
can access files on the cluster using a separate data channel with
privileged access. By default, it is ``_no_`` .

``[-privileged-user-name <text>]`` - User Name for Privileged Access::
This parameter specifies the privileged user name. It is used to specify
the privileged user name for accessing files on the cluster using a
separate data channel with privileged access. The input for this field
should be in ``_"domain\user name"_`` format. If `-allow-privileged-
access` is set to ``_no_`` , any value set for this field is ignored.

``[-is-passthrough-read-enabled {true|false}]`` - Is Passthrough Read
Enabled::
This parameter specifies whether passthrough-read should be allowed for
FPolicy servers registered for the policy. Passthrough-read is a way to
read data for offline files without restoring the files to primary
storage. Offline files are the files which have been moved to secondary
storage. If passthrough-read is enabled, the FPolicy server provides the
data for the file over a separate channel instead of restoring the file to
primary storage. By default, this parameter is ``_false_`` .

``[-persistent-store <text>]`` - Persistent Store Name::
This parameter specifies persistent storage name. This can then be used
for enabling the Persistent mode for Fpolicy events.

== Examples

The following example modifies an FPolicy policy.
```

cluster1::> vserver fpolicy policy modify -vserver vs1.example.com -policy-name vs1_pol -events cserver_evt,v1e1
-engine native -is-mandatory true -allow-privileged-access no -is-passthrough-read-enabled false

cluster1::> vserver fpolicy policy show -vserver vs1.example.com -policy-name vs1_pol
Vserver: vs1.example.com
Policy Name: vs1_pol
Events to Monitor: cserver_evt, v1e1
FPolicy Engine: native
Is Mandatory Screening Required: true

Allow Privileged Access: no
User Name for Privileged Access: -
Is Passthrough Read Enabled: false
persistent-store: -

```
[[ID456b02645fccc938fffbf4ca9e79fc89]]
= vserver fpolicy policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display policy configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy show` command displays information about all
FPolicy policies belonging to the Vserver. Any Vserver administrator can
see FPolicy policies associated with their Vserver as well as policies
created by the cluster administrator. The command output depends on the
parameter or parameters specified with the command. If you do not specify
any parameters, the command displays the following information about all
FPolicy policies:

* Vserver name
* Policy name
* Events to monitor
* FPolicy engine
* Is mandatory screening required
* Allow privileged access
* User name for privileged access



You can specify the `-fields` parameter to specify which fields of
information to display about FPolicy policies. You can specify additional
parameters to display only information that matches those parameters. For
example, to display information only about FPolicy policies where the
FPolicy server requires privileged access, run the command with the `-
fields` parameter set to policy-name (no "-") and `-allow-privileged-
access` parameter set to ``_yes_`` .

You can specify the `-instance` parameter to display all information for
```

```
all policies in the list form.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the FPolicy policies for the specified Vserver. FPolicy policies created
by the cluster administrator are visible for all Vservers.

``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
the FPolicy policy that you specify.

``[-events <Event name>,...]`` - Events to Monitor::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified event or events.

``[-engine <Engine name>]`` - FPolicy Engine::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified engine.

``[-is-mandatory {true|false}]`` - Is Mandatory Screening Required::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified mandatory attribute.

``[-allow-privileged-access {yes|no}]`` - Allow Privileged Access::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified privileged access.

``[-privileged-user-name <text>]`` - User Name for Privileged Access::
If you specify this parameter, the command displays information only about
the FPolicy policy or policies that use the specified privileged user
name.

``[-is-passthrough-read-enabled {true|false}]`` - Is Passthrough Read
Enabled::
If you specify this parameter, the command displays information only about
the FPolicy policies that use the specified passthrough-read setting.
```

```
``[-persistent-store <text>]`` - Persistent Store Name::
This parameter specifies persistent storage name. This can then be used
for enabling the Persistent mode for Fpolicy events.

== Examples

The following example displays the information about FPolicy policies on
the cluster using the `vserver fpolicy policy show` command.
```

cluster1::> vserver fpolicy policy show
Vserver Policy Events Engine Is Mandatory PrivAccess
--- --- --- --- --- ---
Cluster cserver_pol cserver_ cserver_eng true yes
evt
vs1.example.com p r n true no
vs1.example.com cserver_pol cserver_ cserver_eng true yes
evt
vs2.example.com cserver_pol cserver_ cserver_eng true yes
evt
4 entries were displayed.

```
The following example displays FPolicy policy name information about all
Vserver FPolicy policies with the `-allow-privileged-access` parameter set
to "yes".
```

cluster1::> vserver fpolicy policy show -fields policy-name -allow-privileged-access yes
vserver policy-name
--- ---
Cluster cserver_pol
vs1.example.com cserver_pol
vs2.example.com cserver_pol
3 entries were displayed.

```
[[ID3777776218b1f4c90a3fe387f2042029]]
= vserver fpolicy policy event create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an event

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `vserver fpolicy policy event create` command creates an FPolicy
event. An event describes what to monitor. An event can contain protocol,
file operations, filters, and volume operation event types. In the FPolicy
configuration, an event is attached to an FPolicy policy. You can attach
the same event to one or more policies.

NOTE: Three parameters have dependency rules: `-protocol` , `-files-
operations` and `-filters` . The following combinations are supported:

*  Both `-protocol` and `-file-operations`
*  All of `-protocol` , `-file-operations` and `-filters`
*  Specify none of three





== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
create an FPolicy event.

``-event-name <Event name>`` - Event::
This parameter specifies the name of the FPolicy event that you want to
create. An event name can be up to 256 characters long. An event name
value is a string that can only contain any combination of ASCII-range
alphanumeric characters (a-z, A-Z, 0-9), "_" and ".".

``[-protocol <Protocol>]`` - Protocol::
This parameter specifies the protocol name for which the event will be
created. By default, no protocol is selected. The value of this parameter
must be one of the following:
+
  * ``_ cifs _`` - This specifies that the event is for the CIFS protocol.
  * ``_ nfsv3_`` - This specifies that the event is for the NFSv3
protocol.
  * ``_ nfsv4_`` - This specifies that the event is for the NFSv4
protocol.



+
+
NOTE: If you specify `-protocol` , then you must also specify a valid
value for the `-file-operations` parameter.
+
```

```
``[-file-operations <File Operation>,...]`` - File Operations::
This parameter specifies a list of file operations for the FPolicy event.
The event will check the operations specified in this list from all client
requests using the protocol specified in the `-protocol` parameter. The
list can include one or more of the following operations:
+
  * ``_close _`` - File close operations.
  * ``_create _`` - File create operations.
  * ``_create_dir _`` - Directory create operations.
  * ``_delete _`` - File delete operations.
  * ``_delete_dir _`` - Directory delete operations.
  * ``_getattr _`` - Get attribute operations.
  * ``_link _`` - Link operations.
  * ``_lookup _`` - Lookup operations.
  * ``_open _`` - File open operations.
  * ``_read _`` - File read operations.
  * ``_write _`` - File write operations.
  * ``_rename _`` - File rename operations.
  * ``_rename_dir _`` - Directory rename operations.
  * ``_setattr _`` - Set attribute operations.
  * ``_symlink _`` - Symbolic link operations.


+
+
NOTE: If you specify `-file-operations` then you must specify a valid
protocol in the `-protocol` parameter.
+


``[-filters <Filter>,...]`` - Filters::
This parameter specifies a list of filters of given file operation or
operations for the protocol specified in the `-protocol` parameter. The
values in the `-filters` parameter are used to filter client requests. The
list can include one or more of the following:
+
  * ``_ monitor-ads _`` - Filter the client request for alternate data
stream.
  * ``_ close-with-modification _`` - Filter the client request for close
with modification.
  * ``_ close-without-modification_`` - Filter the client request for
close without modification.
  * ``_ close-with-read _`` - Filter the client request for close with
read.
```

```
  * ``_ first-read _`` - Filter the client requests for the first-read.
When this filter is used for CIFS events, the first-read request within a
CIFS session results in FPolicy processing. When this filter is used for
NFS events, the `-file-session-io-grouping-count` and `-file-session-io-
grouping-duration` configurations determine the first read-request for
which FPolicy processing is done.
  * ``_ first-write _`` - Filter the client requests for the first-write.
When this filter is used for CIFS events, the first-write request within a
CIFS session results in FPolicy processing. When this filter is used for
NFS events, the `-file-session-io-grouping-count` and `-file-session-io-
grouping-duration` configurations determine the first-write request for
which FPolicy processing is done.
  * ``_ offline-bit _`` - Filter the client request for offline bit set.
Setting this filter, FPolicy server receives notification only when
offline files are accessed.
  * ``_ open-with-delete-intent _`` - Filter the client request for open
with delete intent. Setting this filter, FPolicy server receives
notification only when an attempt is made to open a file with the intent
to delete it. This is used by file systems when the FILE_DELETE_ON_CLOSE
flag is specified.
  * ``_ open-with-write-intent _`` - Filter the client request for open
with write intent. Setting this filter, FPolicy server receives
notification only when an attempt is made to open a file with the intent
to write something in it.
  * ``_ write-with-size-change _`` - Filter the client request for write
with size change.
  * ``_ setattr-with-owner-change _`` - Filter the client setattr requests
for changing owner of a file or directory.
  * ``_ setattr-with-group-change _`` - Filter the client setattr requests
for changing group of a file or directory.
  * ``_ setattr-with-sacl-change _`` - Filter the client setattr requests
for changing sacl on a file or directory.
  * ``_ setattr-with-dacl-change _`` - Filter the client setattr requests
for changing dacl on a file or directory.
  * ``_ setattr-with-modify-time-change _`` - Filter the client setattr
requests for changing the modification time of a file or directory.
  * ``_ setattr-with-access-time-change _`` - Filter the client setattr
requests for changing the access time of a file or directory.
  * ``_ setattr-with-creation-time-change _`` - Filter the client setattr
requests for changing the creation time of a file or directory.
  * ``_ setattr-with-mode-change _`` - Filter the client setattr requests
for changing the mode bits on a file or directory.
  * ``_ setattr-with-size-change _`` - Filter the client setattr requests
for changing the size of a file.
  * ``_ setattr-with-allocation-size-change _`` - Filter the client
setattr requests for changing the allocation size of a file.
```

```
  * ``_ exclude-directory _`` - Filter the client requests for directory
operations. When this filter is specified directory operations are not
monitored.


+
+
NOTE: If you specify a value for the `-filters` parameter, then you must
also specify valid values for the `-file-operations` and `-protocol`
parameters.
+

+
NOTE: If the client sends multiple read/write requests simultaneously for
the same file, then the first-read and first-write filters can result in
more than one FPolicy notification.
+


``[-volume-operation {true|false}]`` - Send Volume Operation
Notifications::
This parameter specifies whether volume operations generate notifications
for the FPolicy event. If this field is set to ``_true_`` then FPolicy
sends notifications when volumes are mounted or unmounted. By default, it
is ``_false_`` .

``[-monitor-fileop-failure {true|false}]`` - Send Failed File Operation
Notifications::
This parameter specifies whether failed file operations generate
notifications for the FPolicy event. If field is set to ``_true_`` then
FPolicy sends notifications when the file operations fail due to lack of
permissions. By default, it is ``_false_`` .

== Examples

The following example creates an FPolicy event.
```

cluster1::> vserver fpolicy policy event create -vserver vs1.example.com -event-name cifs_event -protocol cifs -file-operations open,close,read,write -filters first-read,offline-bit
-volume-operation true -monitor-fileop-failure false
cluster1::> vserver fpolicy policy event show -vserver vs1.example.com -event-name cifs_event
Vserver: vs1.example.com
Event Name: cifs_event
Protocol: cifs
File Operations: open, close, read, write
Filters: first-read, offline-bit
Send Volume Operation Notifications: true

Send Failed File Operation Notifications: false

```
The following is a list of supported `-file-operations` and `-filters` for
the ``_CIFS_`` protocol.
```

```
 Supported |
      File |
Operations | Supported Filters
=========================================================================
============
```

```
close      : monitor-ads, close-with-modification, close-without-
modification,
             offline-bit, close-with-read, exclude-directory
create     : monitor-ads, offline-bit
create_dir : none
delete     : monitor-ads, offline-bit
delete_dir : none
getattr    : offline-bit, exclude-directory
open       : monitor-ads, offline-bit, open-with-delete-intent, open-with-
write-intent,
             exclude-directory
read       : monitor-ads, first-read, offline-bit
write      : monitor-ads, first-write, offline-bit, write-with-size-change
rename     : offline-bit, monitor-ads
rename_dir : none
setattr    : offline-bit, monitor-ads, setattr-with-owner-change,
             setattr-with-group-change, setattr-with-sacl-change,
             setattr-with-dacl-change, setattr-with-modify-time-change,
             setattr-with-access-time-change, setattr-with-creation-time-
change,
             setattr-with-size-change, setattr-with-allocation-size-
change,
             exclude-directory
```

```
The following is a list of supported `-file-operations` and `-filters` for
the ``_nfsv3_`` protocol.
```

```
 Supported |
      File |
Operations | Supported Filters
==========================================================================
============
 create     : offline-bit
 create_dir : none
 delete     : offline-bit
 delete_dir : none
 link       : offline-bit
 lookup     : offline-bit, exclude-directory
 read       : offline-bit, first-read
 write      : offline-bit, write-with-size-change, first-write
 rename     : offline-bit
 rename_dir : none
 setattr    : offline-bit, setattr-with-owner-change, setattr-with-group-
change,
              setattr-with-modify-time-change, setattr-with-access-time-
change,
              setattr-with-mode-change, setattr-with-size-change, exclude-
directory
 symlink    : offline-bit
```

```
The following is a list of supported `-file-operations` and `-filters` for
the ``_nfsv4_`` protocol.
```

```
 Supported |
      File |
Operations | Supported Filters
==========================================================================
============
 close      : offline-bit, exclude-directory
 create     : offline-bit
 create_dir : none
 delete     : offline-bit
 delete_dir : none
 getattr    : offline-bit, exclude-directory
 link       : offline-bit
 lookup     : offline-bit, exclude-directory
 open       : offline-bit, exclude-directory
 read       : offline-bit, first-read
 write      : offline-bit, write-with-size-change, first-write
 rename     : offline-bit
 rename_dir : none
 setattr    : offline-bit, setattr-with-owner-change, setattr-with-group-
change,
              setattr-with-sacl-change, setattr-with-dacl-change,
              setattr-with-modify-time-change, setattr-with-access-time-
change,
              setattr-with-size-change, exclude-directory
 symlink    : offline-bit
```

```
The following is a list of supported `-file-operations` for supported
protocol when `-monitor-fileop-failure` is set to true.
```

```
Protocol  | Supported File Operations
==========================================================================
===============
 cifs     : open
 nfsv3    : create, create_dir, read, write, delete, delete_dir, rename,
rename_dir,
            setattr, link
 nfsv4    : open, create, create_dir, read, write, delete, delete_dir,
rename, rename_dir,
            setattr, link
```

```
[[ID059e155d4ecfdfb6e68cc705bb09b9d4]]
= vserver fpolicy policy event delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an event

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy event delete` command deletes an FPolicy
event.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver from which you want to delete an
FPolicy event.

``-event-name <Event name>`` - Event::
This parameter specifies the name of the FPolicy event you want to delete.

== Examples

The following example deletes an FPolicy event.
```

cluster1::> vserver fpolicy policy event delete -vserver vs1.example.com -event-name cifs_event

```
[[ID24a1220bd5d3ee62949cf2e3b12ee4d2]]
= vserver fpolicy policy event modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an event

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver fpolicy policy event modify` command modifies an FPolicy
event. An event describes what to monitor. An event can contain protocol,
file operations, filters, and volume operation event types. In the FPolicy
configuration, an event is attached to an FPolicy policy. You can attach
the same event to one or more policies. You can modify an event while it
is attached to an FPolicy policy. Any changes to the event take effect
immediately.

NOTE: Three parameters have dependency rules: `-protocol` , `-files-
operations` and `-filters` . The following combinations are supported:

*  Both `-protocol` and `-file-operations`
*  All of `-protocol` , `-file-operations` and `-filters`
*  Specify none of three





== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
modify an FPolicy event.

``-event-name <Event name>`` - Event::
This parameter specifies the name of the FPolicy event that you want to
modify. An event name can be up to 256 characters long. An event name
value is a string that can only contain any combination of ASCII-range
alphanumeric characters (a-z, A-Z, 0-9), "_" and "." .

``[-protocol <Protocol>]`` - Protocol::
This parameter specifies the protocol name for which the event will be
modified. By default, no protocol is selected. The value of this parameter
must be one of the following:
+
  * ``_ cifs _`` - This specifies that the event is for the CIFS protocol.
  * ``_ nfsv3_`` - This specifies that the event is for the NFSv3
protocol.
  * ``_ nfsv4_`` - This specifies that the event is for the NFSv4
protocol.


+
```

```
+
NOTE: If you specify `-protocol` , then you must also specify a valid
value for the `-file-operations` parameter.
+


``[-file-operations <File Operation>,...]`` - File Operations::
This parameter specifies a list of file operations for the FPolicy event.
The event will check the operations specified in this list from all client
requests using the protocol specified in the `-protocol` parameter. The
list can include one or more of the following operations:
+
  * ``_ close _`` - File close operations.
  * ``_ create _`` - File create operations.
  * ``_ create_dir_`` - Directory create operations.
  * ``_ delete _`` - File delete operations.
  * ``_ delete_dir_`` - Directory delete operations.
  * ``_ getattr _`` - Get attribute operations.
  * ``_ link _`` - Link operations.
  * ``_ lookup _`` - Lookup operations.
  * ``_ open _`` - File open operations.
  * ``_ read _`` - File read operations.
  * ``_ write _`` - File write operations.
  * ``_ rename _`` - File rename operations.
  * ``_ rename_dir_`` - Directory rename operations.
  * ``_ setattr _`` - Set attribute operations.
  * ``_ symlink _`` - Symbolic link operations.


+
+
NOTE: If you specify `-file-operations` then you must specify a valid
protocol in the `-protocol` parameter.
+


``[-filters <Filter>,...]`` - Filters::
This parameter specifies a list of filters of given file operation or
operations for the protocol specified in the `-protocol` parameter. The
values in the `-filters` parameter are used to filter client requests. The
list can include one or more of the following:
+
  * ``_ monitor-ads _`` - Filter the client request for alternate data
stream.
  * ``_ close-with-modification _`` - Filter the client request for close
with modification.
```

```
  * ``_ close-without-modification_`` - Filter the client request for
close without modification.
  * ``_ close-with-read _`` - Filter the client request for close with
read.
  * ``_ first-read _`` - Filter the client requests for the first-read.
When this filter is used for CIFS events, the first-read request within a
CIFS session results in FPolicy processing. When this filter is used for
NFS events, the `-file-session-io-grouping-count` and `-file-session-io-
grouping-duration` configurations determine the first read-request for
which FPolicy processing is done.
  * ``_ first-write _`` - Filter the client requests for the first-write.
When this filter is used for CIFS events, the first-write request within a
CIFS session results in FPolicy processing. When this filter is used for
NFS events, the `-file-session-io-grouping-count` and `-file-session-io-
grouping-duration` configurations determine the first-write request for
which FPolicy processing is done.
  * ``_ offline-bit _`` - Filter the client request for offline bit set.
Setting this filter, FPolicy server receives notification only when
offline files are accessed.
  * ``_ open-with-delete-intent _`` - Filter the client request for open
with delete intent. Setting this filter, FPolicy server receives
notification only when an attempt is made to open a file with the intent
to delete it. This is used by file systems when the FILE_DELETE_ON_CLOSE
flag is specified.
  * ``_ open-with-write-intent _`` - Filter the client request for open
with write intent. Setting this filter, FPolicy server receives
notification only when an attempt is made to open a file with the intent
to write something in it.
  * ``_ write-with-size-change _`` - Filter the client request for write
with size change.
  * ``_ setattr-with-owner-change _`` - Filter the client setattr requests
for changing owner of a file or directory.
  * ``_ setattr-with-group-change _`` - Filter the client setattr requests
for changing group of a file or directory.
  * ``_ setattr-with-sacl-change _`` - Filter the client setattr requests
for changing sacl on a file or directory.
  * ``_ setattr-with-dacl-change _`` - Filter the client setattr requests
for changing dacl on a file or directory.
  * ``_ setattr-with-modify-time-change _`` - Filter the client setattr
requests for changing the modification time of a file or directory.
  * ``_ setattr-with-access-time-change _`` - Filter the client setattr
requests for changing the access time of a file or directory.
  * ``_ setattr-with-creation-time-change _`` - Filter the client setattr
requests for changing the creation time of a file or directory.
  * ``_ setattr-with-mode-change _`` - Filter the client setattr requests
for changing the mode bits on a file or directory.
```

```
  * ``_ setattr-with-size-change _`` - Filter the client setattr requests
for changing the size of a file.
  * ``_ setattr-with-allocation-size-change _`` - Filter the client
setattr requests for changing the allocation size of a file.
  * ``_ exclude-directory _`` - Filter the client requests for directory
operations. When this filter is specified directory operations are not
monitored.


+
+
NOTE: If you specify a value for the `-filters` parameter, then you must
also specify valid values for the `-file-operations` and `-protocol`
parameters.
+

+
NOTE: If the client sends multiple read/write requests simultaneously for
the same file, then the first-read and first-write filters can result in
more than one FPolicy notification.
+


``[-volume-operation {true|false}]`` - Send Volume Operation
Notifications::
This parameter specifies whether volume operations generate notifications
for the FPolicy event. If this field is set to ``_true_`` then FPolicy
sends notifications when volumes are mounted or unmounted. By default, it
is ``_false_`` .

``[-monitor-fileop-failure {true|false}]`` - Send Failed File Operation
Notifications::
This parameter specifies whether failed file operation generate
notifications for the FPolicy event. If field is set to ``_true_`` then
FPolicy sends notifications when the file operations fail due to lack of
permissions. By default, it is ``_false_`` .

== Examples

The following example modifies an FPolicy event.
```

cluster1::> vserver fpolicy policy event modify -vserver vs1.example.com -event-name cifs_event -protocol cifs -file-operations open,close,read,write -filters first-read,offline-bit
-volume-operation true -monitor-fileop-failure false
cluster1::> vserver fpolicy policy event show -vserver vs1.example.com -event-name cifs_event
Vserver: vs1.example.com
Event Name: cifs_event

Protocol: cifs
File Operations: open, close, read, write
Filters: first-read, offline-bit
Send Volume Operation Notifications: true
Send Failed File Operation Notifications: false

```
The following is a list of supported `-file-operations` and `-filters` for
the ``_CIFS_`` protocol.
```

```
 Supported |
      File |
Operations | Supported Filters
==============================================================================
============
```

```
close      : monitor-ads, close-with-modification, close-without-
modification,
             offline-bit, close-with-read, exclude-directory
create     : monitor-ads, offline-bit
create_dir : none
delete     : monitor-ads, offline-bit
delete_dir : none
getattr    : offline-bit, exclude-directory
open       : monitor-ads, offline-bit, open-with-delete-intent, open-with-
write-intent,
             exclude-directory
read       : monitor-ads, first-read, offline-bit
write      : monitor-ads, first-write, offline-bit, write-with-size-change
rename     : offline-bit, monitor-ads
rename_dir : none
setattr    : offline-bit, monitor-ads, setattr-with-owner-change,
             setattr-with-group-change, setattr-with-sacl-change,
             setattr-with-dacl-change, setattr-with-modify-time-change,
             setattr-with-access-time-change, setattr-with-creation-time-
change,
             setattr-with-size-change, setattr-with-allocation-size-
change,
             exclude-directory
```

```
The following is a list of supported `-file-operations` and `-filters` for
the ``_nfsv3_`` protocol.
```

```
 Supported |
      File |
Operations | Supported Filters
==========================================================================
============
 create     : offline-bit
 create_dir : none
 delete     : offline-bit
 delete_dir : none
 link       : offline-bit
 lookup     : offline-bit, exclude-directory
 read       : offline-bit, first-read
 write      : offline-bit, write-with-size-change, first-write
 rename     : offline-bit
 rename_dir : none
 setattr    : offline-bit, setattr-with-owner-change, setattr-with-group-
change,
              setattr-with-modify-time-change, setattr-with-access-time-
change,
              setattr-with-mode-change, setattr-with-size-change, exclude-
directory
 symlink    : offline-bit
```

```
The following is a list of supported `-file-operations` and `-filters` for
the ``_nfsv4_`` protocol.
```

```
 Supported |
      File |
Operations | Supported Filters
==========================================================================
============
 close      : offline-bit, exclude-directory
 create     : offline-bit
 create_dir : none
 delete     : offline-bit
 delete_dir : none
 getattr    : offline-bit, exclude-directory
 link       : offline-bit
 lookup     : offline-bit, exclude-directory
 open       : offline-bit, exclude-directory
 read       : offline-bit, first-read
 write      : offline-bit, write-with-size-change, first-write
 rename     : offline-bit
 rename_dir : none
 setattr    : offline-bit, setattr-with-owner-change, setattr-with-group-
change,
              setattr-with-sacl-change, setattr-with-dacl-change,
              setattr-with-modify-time-change, setattr-with-access-time-
change,
              setattr-with-size-change, exclude-directory
 symlink    : offline-bit
```

```
The following is a list of supported `-file-operations` for supported
protocol when `-monitor-fileop-failure` is set to true.
```

```
Protocol  | Supported File Operations
==========================================================================
===============
 cifs     : open
 nfsv3    : create, create_dir, read, write, delete, delete_dir, rename,
rename_dir,
            setattr, link
 nfsv4    : open, create, create_dir, read, write, delete, delete_dir,
rename, rename_dir,
            setattr, link
```

```
[[ID18f796a7c3776ea7cfb2de23a5578407]]
= vserver fpolicy policy event show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display events

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy event show` command displays information about
all FPolicy events belonging to the Vserver. Any Vserver administrator can
see FPolicy events associated with their Vserver as well as FPolicy events
created by the cluster administrator. The command output depends on the
parameter or parameters specified with the command. If you do not specify
any parameters, the command displays the following information about all
FPolicy events:

* Vserver name
* FPolicy event name
* Protocol name
* List of file operations
* List of filters
* Volume operation
* Monitor failed file operation



You can specify the `-fields` parameter to specify which fields of
information to display about FPolicy events. You can specify additional
parameters to display only information that matches those parameters. For
example, to display information only about all CIFS events configured with
the `-volume-operation` field set, run the command with the `-fields`
parameter set to `-event-name` event-name `-protocol```_cifs_```-volume-
operation```_yes_`` .

You can specify the `-instance` parameter to display all information for
all policies in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.
```

```
| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the FPolicy events for the specified Vserver. Events created on the admin
Vserver by the cluster administrator are visible in all Vservers.

``[-event-name <Event name>]`` - Event::
If you specify this parameter, the command displays information only about
the FPolicy event that matches the specified event name.

``[-protocol <Protocol>]`` - Protocol::
If you specify this parameter, the command displays information only about
the FPolicy event or events that use the specified protocol.

``[-file-operations <File Operation>,...]`` - File Operations::
If you specify this parameter, the command displays information only about
the FPolicy event or events that use the specified file operation or
operations.

``[-filters <Filter>,...]`` - Filters::
If you specify this parameter, the command displays information only about
the FPolicy event or events that use the specified filter or filters.

``[-volume-operation {true|false}]`` - Send Volume Operation
Notifications::
If this field is set to ``_true_`` , then FPolicy displays information
about those events for which it sends notifications when volumes are
mounted or unmounted. If you set this parameter to ``_true_`` , the
command displays information about events where the `-volume-operation`
parameter is set ``_true_`` and volume operations such as mount and
unmount are monitored. If you set this parameter to ``_false_`` , the
command displays information about events where volume operations are not
monitored.

``[-monitor-fileop-failure {true|false}]`` - Send Failed File Operation
Notifications::
If you specify this parameter, the command displays information only about
the FPolicy event or events that has `-monitor-fileop-failure` parameter
set to the specified value.

== Examples
```

```
The following example displays the information about all Vserver FPolicy
policy events.
```

cluster1::> vserver fpolicy policy event show
Event File Volume
Vserver Name Protocols Operations Filters Operation
--- --- --- --- --- ---
Cluster cserver_evt cifs open, close, first-write, true
read, write first-read
vs1.example.com cserver_evt cifs open, close, first-write, true
read, write first-read
vs1.example.com v1e1 cifs open, read first-read -
vs1.example.com v1e2 cifs open - false
vs1.example.com v1e3 nfsv4 open - true
vs2.example.com cserver_evt cifs open, close, first-write, true
read, write first-read
6 entries were displayed.

```
The following example displays event name information about all Vserver
FPolicy policy events with CIFS as a protocol and with false as volume
operation.
```

cluster1::> vserver fpolicy policy event show -fields event-name -protocol cifs -volume-operation false
vserver event-name
--- ---
vs1.example.com v1e2

```
[[ID99e100a76c7c7d165409e7377c63ae95]]
= vserver fpolicy policy external-engine create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an external engine

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy external-engine create` command creates an
FPolicy external engine. The cluster uses the external engine to hold
configuration information that it needs in order to send notification
information to the FPolicy servers. It specifies the primary servers and
secondary servers to which the cluster will send notifications. It also
```

```
specifies FPolicy server related configuration information.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
create an FPolicy external engine.

``-engine-name <Engine name>`` - Engine::
This parameter specifies the name of the FPolicy external engine that you
want to create. An external engine name can be up to 256 characters long.
An external engine name is a string that can only contain any combination
of ASCII-range alphanumeric characters (a-z, A-Z, 0-9), "_", and "." .

``-primary-servers <IP Address>,...`` - Primary FPolicy Servers::
This parameter specifies a list of IP addresses for the primary FPolicy
servers to which you want the external engine you create to apply. The `-
primary-servers` parameter is used to specify a list of servers to which
to send file access events for a given FPolicy policy. When an
administrator configures multiple servers as primary servers,
notifications are sent to the FPolicy servers in a round-robin fashion.

``-port <integer>`` - Port Number of FPolicy Service::
This parameter specifies the port number for the FPolicy service.

``[-secondary-servers <IP Address>,...]`` - Secondary FPolicy Servers::
This parameter specifies a list of IP addresses for the secondary FPolicy
servers to which you want the external engine you create to apply.
Secondary servers will be used only when all the primary servers are not
reachable. When an administrator configures multiple servers as secondary
servers, notifications are sent to FPolicy server in a round-robin
fashion. By default, no secondary server is selected.

``[-extern-engine-type <External Engine Type>]`` - External Engine Type::
This parameter specifies the type of the external engine. This specifies
how the FPolicy server should behave, synchronously or asynchronously. By
default, it is ``_synchronous_`` in nature. When set to ``_synchronous_``
, after sending a notification to the external FPolicy server, request
processing does not continue until after receiving a response from the
FPolicy server. At that point request flow either continues or processing
results in denial, depending on whether the response from the FPolicy
server permits the requested action. When set to ``_asynchronous_`` ,
after sending a notification to the external FPolicy server, file request
processing continues.

``-ssl-option {no-auth|server-auth|mutual-auth}`` - SSL Option for
```

```
External Communication::
This parameter specifies the SSL option for external communication with
the FPolicy server. Possible values include the following:
+
  * no-auth : When set to no-auth, no authentication takes place. The
communication link is established over the TCP protocol.
  * server-auth : When set to server-auth, only the FPolicy server is
authenticated by the Vserver. With this option, before creating the
FPolicy external engine, the administrator must install the public
certificate of the certificate authority (CA) that signed the FPolicy
server certificate.
  * mutual-auth : When set to mutual-auth, mutual authentication takes
place between the Vserver and the FPolicy server, i.e. authentication of
the FPolicy server by the Vserver along with authentication of the Vserver
by the FPolicy server. With this option, before creating the FPolicy
external engine, the administrator must install the public certificate of
the certificate authority (CA) that signed the FPolicy server certificate
along with the public certificate and key file for authentication of the
Vserver.



+
The public certificate of certificate authority (CA) that is used to sign
the FPolicy server certificate is installed using the
xref:{relative_path}security-certificate-install.html[security certificate
install] command with `-type` set to ``_client_ca_`` . The private key and
public certificate required for authentication of the Vserver is installed
using the xref:{relative_path}security-certificate-install.html[security
certificate install] command with `-type` set to ``_server_`` .

``[-reqs-cancel-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Canceling a Request [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout for canceling a request. It is used
to specify the time interval in which the node waits for a response from
the FPolicy server. Beyond this timeout, a cancel request is sent to the
FPolicy server to cancel the pending request. The request is then sent to
an alternate FPolicy server that is registered for the policy. This
timeout helps in handling a FPolicy server that is not responding, which
can improve CIFS/NFS client response. Also, this feature can help in
releasing of system resources since the request is moved from a down/bad
FPolicy server to an alternate FPolicy server. The value for this field
must be between 0s and 100s. By default, it is 20s.

``[-reqs-abort-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
```

```
Aborting a Request [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout for aborting a request. The value for
this field must be between 0s and 200s. By default, it is 40s.

``[-status-req-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Interval for
Sending Status Requests [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the interval for sending status requests. It is
used to specify the interval after which a status request will be send to
the FPolicy server. The value for this field must be between 0s and 50s.
By default, it is 10s.

``[-max-connection-retries <integer>]`` - Max Reconnect Attempt [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the maximum number of attempts to reconnect to
the FPolicy server from a Vserver. It is used to specify the number of
times a broken connection will be retried. The value for this field must
be between 0 and 20. By default, it is 5.

``[-max-server-reqs <integer>]`` - Maximum Outstanding Requests for
FPolicy Server [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the maximum number of outstanding requests for
the FPolicy server. It is used to specify maximum outstanding requests
that will be queued up for the FPolicy server. The value for this field
must be between 1 and 10000. The default values are 500, 1000 or 2000 for
Low-end(< 64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB
memory) Platforms respectively.

``[-server-progress-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Disconnecting Non-responsive Server [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the timeout for disconnecting non-responsive
FPolicy servers. It is used to specify the time interval after which the
connection to the FPolicy server is terminated. This happens only when the
FPolicy server's queue contains the maximum allowed number of requests
that it can hold in its queue and no response is received within this
timeout. The maximum allowed number of requests is either 50 (the default)
or the number specified by the `-max-server-reqs` parameter. The value for
this field must be between 1s and 100s. By default, it is 60s.

``[-keep-alive-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Interval for
Sending Keep-Alive Messages [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the interval in hours (h), minutes (m), or
seconds (s) at which keep-alive messages are sent to the FPolicy server.
```

```
Keep-alive messages are used to detect half-open connections. The range of
supported values for this field is 10 through 600 (h, m, or s).
Alternatively, the value can be set to 0, which disables keep-alive
messages and prevents them from being sent to the FPolicy servers. The
default value for this field is 120s.

``[-certificate-common-name <FQDN or Custom Common Name>]`` - FQDN or
Custom Common Name::
This parameter specifies the certificate name as a fully qualified domain
name (FQDN) or custom common name. The certificate is used if SSL
authentication between the Vserver and the FPolicy server is configured.

``[-certificate-serial <text>]`` - Serial Number of Certificate::
This parameter specifies the serial number of the certificate used for
authentication if SSL authentication between the Vserver and the FPolicy
server is configured.

``[-certificate-ca <text>]`` - Certificate Authority::
This parameter specifies the certificate authority (CA) name of the
certificate used for authentication if SSL authentication between the
Vserver and the FPolicy server is configured.

``[-recv-buffer-size <integer>]`` - Receive Buffer Size [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the receive buffer size of the connected socket
for the FPolicy server. The default value is set to 256 kilobytes (Kb).
When the value is set to 0, the size of the receive buffer is set to a
value defined by the system. For example, if the default receive buffer
size of the socket is 65536 bytes, by setting the tunable value to 0, the
socket buffer size is set to 65536 bytes. You can use any non-default
value to set the size (in bytes) of the receive buffer.

``[-send-buffer-size <integer>]`` - Send Buffer Size [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the send buffer size of the connected socket for
the FPolicy server. The default value is set to 1 Mb. When the value is
set to 0, the size of the send buffer is set to a value defined by the
system. For example, if the default send buffer size of the socket is set
to 65536 bytes, by setting the tunable value to 0, the socket buffer size
is set to 65536 bytes. You can use any non-default value to set the size
(in bytes) of the send buffer.

``[-session-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Session ID Purge Timeout During Reconnection [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the interval after which a new session ID is sent
```

```
to the FPolicy server during reconnection attempts. The value for this
field must be between 0s and 200s. The default value is set to 10 seconds.
If the connection between the storage controller and the FPolicy server is
terminated and reconnection is made within the `-session-timeout`
interval, the old session ID is sent to FPolicy server so that it can send
responses for old notifications.

``[-is-resiliency-enabled {true|false}]`` - Is Resiliency Feature
Enabled::
This parameter specifies whether the resiliency feature is enabled. When
this parameter is set to ``_true_`` and all the primary and secondary
servers are down, or no response is received from the FPolicy servers,
file access events are stored inside the storage controller under the
specified `-resiliency-directory-path` . To deny the file access events
from being stored under these circumstances, set this parameter to
``_false_`` . By default, it is ``_false_`` .

``[-resiliency-max-retention-duration
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Maximum
Notification Retention Duration::
This parameter specifies the duration for which the notifications are
written to files inside the storage controller during network outage. The
value for this field must be between 0s and 600s. By default, it is set to
180s.

``[-resiliency-directory-path <text>]`` - Directory for Notification
Storage::
This parameter specifies the directory path under the `-vserver`
namespace, where notifications are stored in the files whenever network
outage happens.

``[-extern-engine-format {xml|protobuf}]`` - External Engine Format::
This parameter specifies the format of the Fpolicy notification messages
sent to the external engine. Valid values: ``_xml_`` or ``_protobuf_`` .
Default value for this parameter is ``_xml_`` . When set to ``_protobuf_``
, the notification messages are encoded in binary form using Google
Protobuf. Before setting this to ``_protobuf_`` , ensure that the Fpolicy
server also supports Protobuf deserialization.

== Examples

The following example creates an FPolicy external engine.
```

cluster1::> vserver fpolicy policy external-engine create -vserver vs1.example.com -engine-name new_engine -primary-servers 1.1.1.1 -port 10 -secondary-servers 2.2.2.2 -ssl-option mutual-auth -extern-engine-type synchronous -extern-engine-format xml -certificate-serial 8DDE112A114D1FBC -certificate-common-name

Sample1-FPolicy-Client -certificate-ca TASample1

```
cluster1::> vserver fpolicy policy external-engine show -vserver vs1.example.com -engine-name new_engine
Vserver: vs1.example.com
Engine: new_engine
Primary FPolicy Servers: 1.1.1.1
Port Number of FPolicy Service: 10
Secondary FPolicy Servers: 2.2.2.2
External Engine Type: synchronous
External Engine Format: xml
SSL Option for External Communication: mutual-auth
FQDN or Custom Common Name: Sample1-FPolicy-Client
Serial Number: 8DDE112A114D1FBC
Certificate Authority: TASample1
```

```
== Related Links

* xref:{relative_path}security-certificate-install.html[security
certificate install]




[[IDa80304723bcf60a7622105b106350fdd]]
= vserver fpolicy policy external-engine delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an external engine

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy external-engine delete` command deletes an
FPolicy external engine.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver from which you want to delete an
FPolicy external engine.

``-engine-name <Engine name>`` - Engine::
This parameter specifies the name of the FPolicy external engine you want
to delete.

== Examples

The following example deletes an FPolicy external engine.
```

cluster1::> vserver fpolicy policy external-engine show -vserver vs1.example.com -engine-name new_engine
Vserver: vs1.example.com
Engine: new_engine
Primary FPolicy Servers: 1.1.1.1
Port Number of FPolicy Service: 10
Secondary FPolicy Servers: 2.2.2.2

External Engine Type: synchronous
SSL Option for External Communication: mutual-auth
FQDN or Custom Common Name: Sample1-FPolicy-Client
Serial Number: 8DDE112A114D1FBC
Certificate Authority: TASample1

cluster1::> vserver fpolicy policy external-engine delete -vserver vs1.example.com -engine-name new_engine

```
[[ID0398d136acee3b8f169cf48fedd53d15]]
= vserver fpolicy policy external-engine modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an external engine

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy external-engine modify` command modifies an
FPolicy external engine. The cluster uses the external engine to hold
configuration information that it needs in order to send notification
information to the FPolicy servers. It specifies the primary servers and
secondary servers to which the cluster will send notifications. It also
specifies FPolicy server related configuration information.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
modify an FPolicy external engine.

``-engine-name <Engine name>`` - Engine::
This parameter specifies the name of the FPolicy external engine that you
want to modify. An external engine name can be up to 256 characters long.
An external engine name is a string that can only contain any combination
of ASCII-range alphanumeric characters (a-z, A-Z, 0-9), "_", and "." .

``[-primary-servers <IP Address>,...]`` - Primary FPolicy Servers::
This parameter specifies a list of IP addresses for the primary FPolicy
servers to which you want the external engine you modify to apply. The `-
primary-servers` parameter is used to specify a list of servers to which
to send file access events for a given FPolicy policy. When an
```

```
administrator configures multiple servers as primary servers,
notifications are sent to the FPolicy servers in a round-robin fashion.

``[-port <integer>]`` - Port Number of FPolicy Service::
This parameter specifies the port number for the FPolicy service.

``[-secondary-servers <IP Address>,...]`` - Secondary FPolicy Servers::
This parameter specifies a list of IP addresses for the secondary FPolicy
servers to which you want the external engine you modify to apply.
Secondary servers will be used only when all the primary servers are not
reachable. When an administrator configures multiple servers as secondary
servers, notifications are sent to FPolicy server in a round-robin
fashion. By default, no secondary server is selected.

``[-extern-engine-type <External Engine Type>]`` - External Engine Type::
This parameter specifies the type of the external engine. This specifies
how the FPolicy server should behave, synchronously or asynchronously. By
default, it is synchronous in nature. When set to synchronous, after
sending a notification to the external FPolicy server, request processing
does not continue until after receiving a response from the FPolicy
server. At that point request flow either continues or processing results
in denial, depending on whether the response from the FPolicy server
permits the requested action. When set to asynchronous, after sending a
notification to the external FPolicy server, file request processing
continues.

``[-ssl-option {no-auth|server-auth|mutual-auth}]`` - SSL Option for
External Communication::
This parameter specifies the SSL option for external communication with
the FPolicy server. Possible values include the following:
+
  * no-auth : When set to no-auth, no authentication takes place. The
communication link is established over the TCP protocol.
  * server-auth : When set to server-auth, only the FPolicy server is
authenticated by the Vserver. With this option, before creating the
FPolicy external engine, the administrator must install the public
certificate of the certificate authority (CA) that signed the FPolicy
server certificate.
  * mutual-auth : When set to mutual-auth, mutual authentication takes
place between the Vserver and the FPolicy server, i.e. authentication of
the FPolicy server by the Vserver along with authentication of the Vserver
by the FPolicy server. With this option, before creating the FPolicy
external engine, the administrator must install the public certificate of
the certificate authority (CA) that signed the FPolicy server certificate
along with the public certificate and key file for authentication of the
Vserver.
```

```
+
The public certificate of certificate authority (CA) that is used to sign
the FPolicy server certificate is installed using the
xref:{relative_path}security-certificate-install.html[security certificate
install] command with `-type` set to ``_client_ca_`` . The private key and
public certificate required for authentication of the Vserver is installed
using the xref:{relative_path}security-certificate-install.html[security
certificate install] command with `-type` set to ``_server_`` .

``[-reqs-cancel-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Canceling a Request [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout for canceling a request. It is used
to specify the time interval in which the node waits for a response from
the FPolicy server. Beyond this timeout, a cancel request is sent to the
FPolicy server to cancel the pending request. The request is then sent to
an alternate FPolicy server that is registered for the policy. This
timeout helps in handling a FPolicy server that is not responding, which
can improve CIFS/NFS client response. Also, this feature can help in
releasing of system resources since the request is moved from a down/bad
FPolicy server to an alternate FPolicy server. The value for this field
must be between 0s and 100s. By default, it is 20s.

``[-reqs-abort-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Aborting a Request [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout for aborting a request. The value for
this field must be between 0s and 200s. By default, it is 40s.

``[-status-req-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Interval for
Sending Status Requests [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the interval for sending status requests. It is
used to specify the interval after which a status request will be send to
the FPolicy server. The value for this field must be between 0s and 50s.
By default, it is 10s.

``[-max-connection-retries <integer>]`` - Max Reconnect Attempt [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the maximum number of attempts to reconnect to
the FPolicy server from a Vserver. It is used to specify the number of
times a broken connection will be retried. The value for this field must
be between 0 and 20. By default, it is 5.

``[-max-server-reqs <integer>]`` - Maximum Outstanding Requests for
```

```
FPolicy Server [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the maximum number of outstanding requests for
the FPolicy server. It is used to specify the maximum outstanding requests
that will be queued up for the FPolicy server. The value for this field
must be between 1 and 10000. The default values are 500, 1000 or 2000 for
Low-end(< 64 GB memory), Mid-end(>=64 GB memory) and High-end(>=128 GB
memory) Platforms respectively.

``[-server-progress-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Disconnecting Non-responsive Server [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the timeout for disconnecting non-responsive
FPolicy servers. It is used to specify the time interval after which the
connection to the FPolicy server is terminated. This happens only when the
FPolicy server's queue contains the maximum allowed number of requests
that it can hold in its queue and no response is received within this
timeout. The maximum allowed number of requests is either 50 (the default)
or the number specified by the `-max-server-reqs` parameter. The value for
this field must be between 1s and 100s. By default, it is 60s.

``[-keep-alive-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Interval for
Sending Keep-Alive Messages [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the interval in hours (h), minutes (m), or
seconds (s) at which keep-alive messages are sent to the FPolicy server.
Keep-alive messages are used to detect half-open connections. The range of
supported values for this field is 10 through 600 (h, m, or s).
Alternatively, the value can be set to 0, which disables keep-alive
messages and prevents them from being sent to the FPolicy servers. The
default value for this field is 120s.

``[-certificate-common-name <FQDN or Custom Common Name>]`` - FQDN or
Custom Common Name::
This parameter specifies the certificate name as a fully qualified domain
name (FQDN) or custom common name. The certificate is used if SSL
authentication between the Vserver and the FPolicy server is configured.

``[-certificate-serial <text>]`` - Serial Number of Certificate::
This parameter specifies the serial number of the certificate used for
authentication if SSL authentication between the Vserver and the FPolicy
server is configured.

``[-certificate-ca <text>]`` - Certificate Authority::
This parameter specifies the certificate authority (CA) name of the
certificate used for authentication if SSL authentication between the
```

```
Vserver and the FPolicy server is configured.

``[-recv-buffer-size <integer>]`` - Receive Buffer Size [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the receive buffer size of the connected socket
for the FPolicy server. The default value is set to 256 kilobytes (Kb).
When the value is set to 0, the size of the receive buffer is set to a
value defined by the system. For example, if the default receive buffer
size of the socket is 65536 bytes, by setting the tunable value to 0, the
socket buffer size is set to 65536 bytes. You can use any non-default
value to set the size (in bytes) of the receive buffer.

``[-send-buffer-size <integer>]`` - Send Buffer Size [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the send buffer size of the connected socket for
the FPolicy server. The default value is set to 1 Mb. When the value is
set to 0, the size of the send buffer is set to a value defined by the
system. For example, if the default send buffer size of the socket is set
to 65536 bytes, by setting the tunable value to 0, the socket buffer size
is set to 65536 bytes. You can use any non-default value to set the size
(in bytes) of the send buffer.

``[-session-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Session ID Purge Timeout During Reconnection [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the interval after which a new session ID is sent
to the FPolicy server during reconnection attempts. The value for this
field must be between 0s and 200s. The default value is set to 10 seconds.
If the connection between the storage controller and the FPolicy server is
terminated and reconnection is made within the `-session-timeout`
interval, the old session ID is sent to FPolicy server so that it can send
responses for old notifications.

``[-is-resiliency-enabled {true|false}]`` - Is Resiliency Feature
Enabled::
This parameter specifies whether the resiliency feature is enabled. When
this parameter is set to ``_true_`` and all the primary and secondary
servers are down, or no response is received from the FPolicy servers,
file access events are stored inside the storage controller under the
specified `-resiliency-directory-path` . To deny the file access events
from being stored under these circumstances, set this parameter to
``_false_`` . By default, it is ``_false_`` .

``[-resiliency-max-retention-duration
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Maximum
Notification Retention Duration::
```

```
This parameter specifies the duration for which the notifications are
written to files inside the storage controller during network outage. The
value for this field must be between 0s and 600s. By default, it is set to
180s.

``[-resiliency-directory-path <text>]`` - Directory for Notification
Storage::
This parameter specifies the directory path under the `-vserver`
namespace, where notifications are stored in the files whenever network
outage happens.

``[-extern-engine-format {xml|protobuf}]`` - External Engine Format::
This parameter specifies the format of the Fpolicy notification messages
sent to the external engine. Valid values: ``_xml_`` or ``_protobuf_`` .
Default value for this parameter is ``_xml_`` . When set to ``_protobuf_``
, the notification messages are encoded in binary form using Google
Protobuf. Before setting this to ``_protobuf_`` , ensure that the Fpolicy
server also supports Protobuf deserialization.

== Examples

The following example modifies an FPolicy external engine.
```

cluster1::> vserver fpolicy policy external-engine modify -vserver vs1.example.com -engine-name new_engine -primary-servers 1.1.1.1 -port 10 -secondary-servers 2.2.2.2

cluster1::> vserver fpolicy policy external-engine show -vserver vs1.example.com -engine-name new_engine
Vserver: vs1.example.com
Engine: new_engine
Primary FPolicy Servers: 1.1.1.1
Port Number of FPolicy Service: 10
Secondary FPolicy Servers: 2.2.2.2
External Engine Type: synchronous
External Engine Format: xml
SSL Option for External Communication: mutual-auth
FQDN or Custom Common Name: Sample1-FPolicy-Client
Serial Number: 8DDE112A114D1FBC
Certificate Authority: TASample1

```
The following example shows how to modify `-recv-buffer-size` and `-send-
buffer-size` to a non-default value of 0.
```

cluster1::*> vserver fpolicy policy external-engine modify -vserver vs1.example.com -engine-name new_engine -recv-buffer-size 0 -send-buffer-size 0

```
== Related Links
```

```
* xref:{relative_path}security-certificate-install.html[security
certificate install]




[[ID62094f230b410a2e7d13abbc32b55c02]]
= vserver fpolicy policy external-engine show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display external engines

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy external-engine show` command displays
information about all FPolicy external engines belonging to the Vserver.
Any Vserver administrator can see FPolicy external engines associated to
their Vserver as well as external engines created by cluster
administrator. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information about all FPolicy external
engines:

* Vserver name
* FPolicy external engine name
* List of primary FPolicy servers
* List of secondary FPolicy servers
* Port number for FPolicy service
* FPolicy external engine type
* FPolicy external engine format



You can specify the `-fields` parameter to specify which fields of
information to display about FPolicy external engines. You can specify
additional parameters to display only information that matches those
parameters. For instance, to display information only about all external
engines where the `-port` parameter is set to ``_9_`` , run the command
with the `-field` parameter set to engine-name and `-port` parameter set
to ``_9_`` .
```

```
You can specify the `-instance` parameter to display all information for
all policies in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the FPolicy external engines for the specified Vserver. FPolicy external
engines that the cluster administrator creates are visible in all
Vservers.

``[-engine-name <Engine name>]`` - Engine::
If you specify this parameter, the command displays information only about
the FPolicy external engine that you specify.

``[-primary-servers <IP Address>,...]`` - Primary FPolicy Servers::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified IP addresses
as primary FPolicy servers.

``[-port <integer>]`` - Port Number of FPolicy Service::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified port for the
FPolicy service.

``[-secondary-servers <IP Address>,...]`` - Secondary FPolicy Servers::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified IP addresses
as secondary FPolicy servers.

``[-extern-engine-type <External Engine Type>]`` - External Engine Type::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified external
engine type.

``[-ssl-option {no-auth|server-auth|mutual-auth}]`` - SSL Option for
External Communication::
If you specify this parameter, the command displays information only about
```

```
the FPolicy external engine or engines that use the specified SSL option.

``[-reqs-cancel-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Canceling a Request [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified timeout for
canceling a request.

``[-reqs-abort-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Aborting a Request [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified timeout for
aborting a request.

``[-status-req-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Interval for
Sending Status Requests [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified interval for
sending status requests.

``[-max-connection-retries <integer>]`` - Max Reconnect Attempt [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified maximum
reconnect attempts.

``[-max-server-reqs <integer>]`` - Maximum Outstanding Requests for
FPolicy Server [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified FPolicy
server maximum outstanding requests.

``[-server-progress-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Timeout for
Disconnecting Non-responsive Server [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified timeout for
disconnecting non-responsive server.

``[-keep-alive-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Interval for
Sending Keep-Alive Messages [.advaced-priv]#(privilege: advanced)#::
```

```
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified keep-alive
interval.

``[-certificate-common-name <FQDN or Custom Common Name>]`` - FQDN or
Custom Common Name::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified certificate
common name.

``[-certificate-serial <text>]`` - Serial Number of Certificate::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified certificate
serial number.

``[-certificate-ca <text>]`` - Certificate Authority::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified certificate
authority name.

``[-recv-buffer-size <integer>]`` - Receive Buffer Size [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified receive
buffer size.

``[-send-buffer-size <integer>]`` - Send Buffer Size [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified send buffer
size.

``[-session-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Session ID Purge Timeout During Reconnection [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified session
timeout.

``[-is-resiliency-enabled {true|false}]`` - Is Resiliency Feature
Enabled::
 If you specify this parameter set to ``_true_`` , the command displays
information only about the FPolicy external engine or engines that has the
resiliency feature enabled.

``[-resiliency-max-retention-duration
```

```
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Maximum
Notification Retention Duration::
 If you specify this parameter, the command displays information only
about the FPolicy external engine or engines that use the specified
network outage duration.

``[-resiliency-directory-path <text>]`` - Directory for Notification
Storage::
 If you specify this parameter, the command displays information only
about the FPolicy external engine or engines that use the specified
directory path.

``[-extern-engine-format {xml|protobuf}]`` - External Engine Format::
If you specify this parameter, the command displays information only about
the FPolicy external engine or engines that use the specified engine
format.

== Examples

The following example displays the information about the configured
external engines using the `vserver fpolicy policy external-engine show`
command.
```

cluster1::> vserver fpolicy policy external-engine show
Primary Secondary External External
Vserver Engine Servers Servers Port Engine Type Engine Format
--- --- --- --- --- --- ---
Cluster cserver_eng 9.9.9.9 - 9 synchronous xml
vs1.example.com cserver_eng 9.9.9.9 - 9 synchronous protobuf
vs1.example.com v1n1 1.1.1.1 2.2.2.2 1 synchronous protobuf
vs2.example.com cserver_eng 9.9.9.9 - 9 synchronous xml
vs2.example.com v2n1 3.3.3.3 5.5.5.5 2 synchronous xml
5 entries were displayed.

```
The following example displays the information about all Vserver FPolicy
external engines with the -port parameter set to 9.
```

cluster1::> vserver fpolicy policy external-engine show -fields engine-name -port 9
vserver engine-name
--- ---
Cluster cserver_eng
vs1.example.com cserver_eng
vs2.example.com cserver_eng
3 entries were displayed.

```
The following example displays the values of all the advanced-level
parameters for the external engine v1n1 in Vserver vs1.example.com.
```

cluster1::*> vserver fpolicy policy external-engine show -vserver vs1.example.com -engine-name v1n1 -instance
(vserver fpolicy policy external-engine show)
Vserver: vs1.example.com
Engine: v1n1
Primary FPolicy Servers: 1.1.1.1
Port Number of FPolicy Service: 1
Secondary FPolicy Servers: 2.2.2.2
External Engine Type: synchronous
External Engine Format: protobuf
SSL Option for External Communication: no-auth
Timeout for Canceling a Request: 20s
Timeout for Aborting a Request: 40s
Interval for Sending Status Requests: 10s
Max Reconnect Attempt: 5
Maximum Outstanding Requests for FPolicy Server: 50
Timeout for Disconnecting Non-responsive Server: 1m
Interval for Sending Keep-Alive Messages: 2m
FQDN or Custom Common Name: -
Serial Number of Certificate: -
Certificate Authority: -
Receive Buffer Size: 0
Send Buffer Size: 0
Session ID Purge Timeout During Reconnection: 10s
Is Resiliency Feature Enabled: true
Maximum Notification Retention Duration: 3m
Directory for Notification Storage: /fpolicy

```
[[ID8501c2f96fe55b6f6a23fdc8a5bcc8f0]]
= vserver fpolicy policy scope create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create scope

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy scope create` command creates an FPolicy scope
for an FPolicy policy. A scope defines the boundaries on which the FPolicy
policy will apply. The Vserver is the basic scope boundary. When you
```

```
create a scope for an FPolicy policy, you must define the FPolicy policy
to which it will apply and you must designate to which Vserver you want to
apply the scope. There are a number of parameters that further restrict
the scope within the specified Vserver. You can restrict the scope by
specifying what to include in the scope. Or you can restrict the scope by
specifying what to exclude from the scope. For example, you can restrict
the scope by specifying which volumes to include using the `-volumes-to-
include` parameter or which volumes to exclude using the `-volumes-to-
exclude` parameter. Once you apply a scope to an enabled policy, policy
event checks get applied to the scope defined by this command.

NOTE:  There are special considerations for the scope for a cluster
FPolicy policy. The cluster FPolicy policy is a policy that the cluster
administrator creates for the admin Vserver. If the cluster administrator
also creates the scope for that cluster FPolicy policy, a Vserver
administrator cannot create a scope for that same policy. However, if the
cluster administrator does not create a scope for the cluster FPolicy
policy, then any Vserver administrator can create the scope for that
cluster policy. In the event that the Vserver administrator creates a
scope for that cluster FPolicy policy, the cluster administrator cannot
subsequently create a cluster scope for that same cluster policy. This is
because the cluster administrator cannot override the scope for the same
cluster policy.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
create an FPolicy policy scope.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the FPolicy policy for which you want
to create the scope.

``[-shares-to-include <Share name>,...]`` - Shares to Include::
This parameter specifies a list of shares for file access monitoring. With
this option, the administrator provides a list of shares, separated by
commas. For file access events relative to the specified shares and file
operations monitored by the FPolicy policy, a notification is generated.
The `-shares-to-include ` parameter can contain regular expressions and
can include metacharacters such as "?" and "*".
+
NOTE: When a share is included in the `-shares-to-include` parameter and
the parent volume of the share is included in the `-volumes-to-exclude`
```

```
parameter, `-volumes-to-exclude` has precedence over `-shares-to-include`
.
+


``[-shares-to-exclude <Share name>,...]`` - Shares to Exclude::
This parameter specifies a list of shares to exclude from file access
monitoring. With this option, the administrator provides a list of shares,
separated by commas. When a share is specified in the `-shares-to-exclude`
parameter, no notification is sent for files accessed relative to that
share. The `-shares-to-exclude` parameter can contain regular expressions
and can include metacharacters such as "?" and "*".

``[-volumes-to-include <volume name>,...]`` - Volumes to Include::
This parameter specifies a list of volumes for file access monitoring.
With this option, the administrator provides a list of volumes, separated
by commas. For file access events within the volume and file operations
monitored by the FPolicy policy, a notification is generated. The `-
volumes-to-include` parameter can contain regular expressions and can
include metacharacters such as "?" and "*".

``[-volumes-to-exclude <volume name>,...]`` - Volumes to Exclude::
This parameter specifies a list of volumes to exclude from file access
monitoring. With this option, the administrator provides a list of
volumes, separated by commas, for which no file access notifications are
generated. The `-volumes-to-exclude` parameter can contain regular
expressions and can include metacharacters such as "?" and "*".
+
NOTE: When a share is included in the `-shares-to-include` parameter and
the parent volume of the share is included in the `-volumes-to-exclude`
parameter, `-volumes-to-exclude` has precedence over `-shares-to-include`
. Similarly, when an export policy is included in the `-export-policies-
to-include` parameter and the parent volume of the export-policy is
included in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has
precedence over `-export-policies-to-include` .
+


``[-export-policies-to-include <FPolicy export policy>,...]`` - Export
Policies to Include::
This parameter specifies a list of export policies for file access
monitoring. With this option, the administrator provides a list of export
policies, separated by commas. For file access events within an export
policy and file operations monitored by the FPolicy policy, a notification
is generated. The `-export-policies-to-include` parameter can contain
regular expressions and can include metacharacters such as "?" and "*".
```

```
+
NOTE: When an export policy is included in the `-export-policies-to-
include` parameter and the parent volume of the export policy is included
in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has
precedence over `-export-policies-to-include` .
+


``[-export-policies-to-exclude <FPolicy export policy>,...]`` - Export
Policies to Exclude::
This parameter specifies a list of export policies to exclude from file
access monitoring. With this option, the administrator provides a list of
export policies, separated by commas, for which no file access
notification is sent. The `-export-policies-exclude` parameter can contain
regular expressions and can include metacharacters such as "?" and *.

``[-file-extensions-to-include <File extension>,...]`` - File Extensions
to Include::
This parameter specifies a list of file extensions, separated by commas,
for a given FPolicy policy for which FPolicy processing is required. Any
file access to files with the same extensions included in the `-file-
extensions-to-include` parameter generates a notification. The `-file-
extensions-to-include` parameter can contain regular expressions and can
include metacharacters such as "?".

``[-file-extensions-to-exclude <File extension>,...]`` - File Extensions
to Exclude::
This parameter specifies a list of file extensions, separated by commas,
for a given FPolicy policy for which FPolicy processing will be excluded.
Using the exclude list, the administrator can request notification for all
extensions except those in the excluded list. Any file access to files
with the same extensions included in the `-file-extensions-to-exclude`
parameter does not generate a notification. The `-file-extensions-to-
exclude` parameter can contain regular expressions and can include
metacharacters such as "?".
+
NOTE: An administrator can specify both `-file-extensions-to-include` and
`-file-extensions-to-exclude` lists. The `-file-extensions-to-exclude`
parameter is checked first before the `-file-extensions-to-include`
parameter is checked.
+


``[-is-file-extension-check-on-directories-enabled {true|false}]`` - Is
File Extension Check on Directories Enabled [.advaced-priv]#(privilege:
advanced)#::
```

```
This parameter specifies whether the file name extension checks apply to
directory objects as well. If this parameter is set to true, the directory
objects are subjected to same extension checks as regular files. If this
parameter is set to false, the directory names are not matched for
extensions and notifications would be sent for directories even if their
name extensions do not match. By default, it is ``_true_`` .

``[-is-monitoring-of-objects-with-no-extension-enabled {true|false}]`` -
Is Monitoring of Objects with No Extension Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether the extension checks apply to objects
with no extension as well. If this parameter is set to true, the objects
with no extension are also monitored along with the objects with
extension. By default, it is ``_false_`` .
+
NOTE: This parameter is ignored when file-extensions-to-include and file-
extensions-to-exclude lists are empty.
+


== Examples

The following example creates an FPolicy policy scope.
```

cluster1::> vserver fpolicy policy scope create -vserver vs1.example.com
-policy-name vs1_pol
-file-extensions-to-include flv,wmv,mp3,mp4
-file-extensions-to-exclude cpp,c,h,txt
cluster1::> vserver fpolicy policy scope show
Vserver Policy Extensions Extensions
Name Name Included Excluded
---------- ---------- ---------- ----------
Cluster cserver_pol txt mp3, wmv
vs1.example.com vs1_pol flv, wmv, mp3, mp4 cpp, c, h, txt
2 entries were displayed.

```
[[ID9d3731c91679e7af0de6be2a1cf6c438]]
= vserver fpolicy policy scope delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete scope

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy scope delete` command deletes an FPolicy
policy scope.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver from which you want to
delete the FPolicy policy scope.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the FPolicy policy for which you want
to delete the scope.

== Examples

The following example deletes a scope of an FPolicy policy.
```

cluster1::> vserver fpolicy policy scope delete -vserver vs1.example.com -policy-name vs1_pol

```
[[ID14de4eaaafbc4712c061eee28706b790]]
= vserver fpolicy policy scope modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify scope

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy scope modify` command modifies an FPolicy
scope for an FPolicy policy. A scope defines the boundaries on which the
FPolicy policy will apply. The Vserver is the basic scope boundary. When
you modify a scope for an FPolicy policy, you must define the FPolicy
policy to which it will apply and you must designate to which Vserver you
want to apply the scope. There are a number of parameters that further
restrict the scope within the specified Vserver. You can restrict the
scope by specifying what to include in the scope. Or you can restrict the
scope by specifying what to exclude from the scope. For example, you can
restrict the scope by specifying which volumes to include using the `-
volumes-to-include` parameter or which volumes to exclude using the `-
volumes-to-exclude` parameter. Once you apply a scope to an enabled
policy, policy event checks get applied to the scope defined by this
command.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
 This parameter specifies the name of the Vserver on which you want to
modify an FPolicy policy scope.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the FPolicy policy for which you want
to modify the scope.

``[-shares-to-include <Share name>,...]`` - Shares to Include::
This parameter specifies a list of shares for file access monitoring. With
this option, the administrator provides a list of shares, separated by
commas. For file access events relative to the specified shares and file
operations monitored by the FPolicy policy, a notification is generated.
The `-shares-to-include` parameter can contain regular expressions and can
include metacharacters such as "?" and "*".
+
NOTE: When a share is included in the `-shares-to-include` parameter and
the parent volume of the share is included in the `-volumes-to-exclude`
parameter, `-volumes-to-exclude` has precedence over `-shares-to-include`
.
+



``[-shares-to-exclude <Share name>,...]`` - Shares to Exclude::
This parameter specifies a list of shares to exclude from file access
```

```
monitoring. With this option, the administrator provides a list of shares,
separated by commas. When a share is specified in the `-shares-to-exclude`
parameter, no notification is sent for files accessed relative to that
share. The `-shares-to-exclude` parameter can contain regular expressions
and can include metacharacters such as "?" and "*".

``[-volumes-to-include <volume name>,...]`` - Volumes to Include::
This parameter specifies a list of volumes for file access monitoring.
With this option, the administrator provides a list of volumes, separated
by commas. For file access events within the volume and file operations
monitored by the FPolicy policy, a notification is generated. The `-
volumes-to-include` parameter can contain regular expressions and can
include metacharacters such as "?" and "*".

``[-volumes-to-exclude <volume name>,...]`` - Volumes to Exclude::
This parameter specifies a list of volumes to exclude from file access
monitoring. With this option, the administrator provides a list of
volumes, separated by commas, for which no file access notifications are
generated. The `-volumes-to-exclude` parameter can contain regular
expressions and can include metacharacters such as "?" and "*".
+
NOTE: When a share is included in the `-shares-to-include` parameter and
the parent volume of the share is included in the `-volumes-to-exclude`
parameter, `-volumes-to-exclude` has precedence over `-shares-to-include`
. Similarly, when an export policy is included in the `-export-policies-
to-include` parameter and the parent volume of the export-policy is
included in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has
precedence over `-export-policies-to-include` .
+


``[-export-policies-to-include <FPolicy export policy>,...]`` - Export
Policies to Include::
This parameter specifies a list of export policies for file access
monitoring. With this option, the administrator provides a list of export
policies, separated by commas. For file access events within an export
policy and file operations monitored by the FPolicy policy, a notification
is generated. The `-export-policies-to-include` parameter can contain
regular expressions and can include metacharacters such as "?" and "*".
+
NOTE: When an export policy is included in the `-export-policies-to-
include` parameter and the parent volume of the export policy is included
in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has
precedence over `-export-policies-to-include` .
+
```

```
``[-export-policies-to-exclude <FPolicy export policy>,...]`` - Export
Policies to Exclude::
This parameter specifies a list of export policies to exclude from file
access monitoring. With this option, the administrator provides a list of
export policies, separated by commas, for which no file access
notification is sent. The `-export-policies-exclude` parameter can contain
regular expressions and can include metacharacters such as "?" and *.

``[-file-extensions-to-include <File extension>,...]`` - File Extensions
to Include::
This parameter specifies a list of file extensions, separated by commas,
for a given FPolicy policy for which FPolicy processing is required. Any
file access to files with the same extensions included in the `-file-
extensions-to-include` parameter generates a notification. The `-file-
extensions-to-include` parameter can contain regular expressions and can
include metacharacters such as "?".

``[-file-extensions-to-exclude <File extension>,...]`` - File Extensions
to Exclude::
This parameter specifies a list of file extensions, separated by commas,
for a given FPolicy policy for which FPolicy processing will be excluded.
Using the exclude list, the administrator can request notification for all
extensions except those in the excluded list. Any file access to files
with the same extensions included in the `-file-extensions-to-exclude`
parameter does not generate a notification. The `-file-extensions-to-
exclude` parameter can contain regular expressions and can include
metacharacters such as "?".
+
NOTE: An administrator can specify both `-file-extensions-to-include` and
`-file-extensions-to-exclude` lists. The `-file-extensions-to-exclude`
parameter is checked first before the `-file-extensions-to-include`
parameter is checked.
+



``[-is-file-extension-check-on-directories-enabled {true|false}]`` - Is
File Extension Check on Directories Enabled [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies whether the file name extension checks apply to
directory objects as well. If this parameter is set to true, the directory
objects are subjected to same extension checks as regular files. If this
parameter is set to false, the directory names are not matched for
extensions and notifications would be sent for directories even if their
name extensions do not match. By default, it is ``_true_`` .

``[-is-monitoring-of-objects-with-no-extension-enabled {true|false}]`` -
```

```
Is Monitoring of Objects with No Extension Enabled [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies whether the extension checks apply to objects
with no extension as well. If this parameter is set to true, the objects
with no extension are also monitored along with the objects with
extension. By default, it is ``_false_`` .
+
NOTE: This parameter is ignored when file-extensions-to-include and file-
extensions-to-exclude lists are empty.
+



== Examples

The following example modifies an FPolicy policy scope.
```

cluster1::> vserver fpolicy policy scope modify -vserver vs1.example.com
-policy-name vs1_pol
-file-extensions-to-include flv,wmv,mp3,mp4
-file-extensions-to-exclude cpp,c,h,txt
cluster1::> vserver fpolicy policy scope show
Vserver Policy Extensions Extensions
Name Name Included Excluded
---------- ---------- ---------- ----------
Cluster cserver_pol txt mp3, wmv
vs1.example.com vs1_pol flv, wmv, mp3, mp4 cpp, c, h, txt
2 entries were displayed.

```
[[ID831122544ff534eaf023cbb1ffbc6d4e]]
= vserver fpolicy policy scope show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display scope

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver fpolicy policy scope show` command displays scope information
about all FPolicy policies belonging to the Vserver. The command output
depends on the parameter or parameters specified with the command. If you
do not specify any parameters, the command displays the following
```

```
information about all FPolicy scopes:

* Vserver name
* Policy name
* The file extensions to include
* The file extensions to exclude



You can use the `-fields` parameter to specify which fields of information
to display about FPolicy scopes. In addition to the fields above, you can
display the following fields:

* The shares to include
* The shares to exclude
* The volumes to include
* The volumes to exclude
* The export policies to include
* The export policies to exclude
* Whether file extension check on directories is enabled
* Whether monitoring of objects with no extension is enabled

You can specify specific parameters to display only information that
matches those parameters. For example, to display scope information only
about all FPolicy policies where the `-file-extensions-to-include`
parameter is set to txt, run the command with the `-fields` parameter set
to policy-name and `-file-extensions-to-include` parameter set to txt.

You can specify the `-instance` parameter to display scope information for
all FPolicy policies in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays scope information only
about the FPolicy policies for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
```

```
the specified FPolicy policy.

``[-shares-to-include <Share name>,...]`` - Shares to Include::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified share or
shares in the include list.

``[-shares-to-exclude <Share name>,...]`` - Shares to Exclude::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified share or
shares in the exclude list.

``[-volumes-to-include <volume name>,...]`` - Volumes to Include::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified volume or
volumes in the include list.

``[-volumes-to-exclude <volume name>,...]`` - Volumes to Exclude::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified volume or
volumes in the exclude list.

``[-export-policies-to-include <FPolicy export policy>,...]`` - Export
Policies to Include::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified export policy
or policies in the include list.

``[-export-policies-to-exclude <FPolicy export policy>,...]`` - Export
Policies to Exclude::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified export policy
or policies in the exclude list.

``[-file-extensions-to-include <File extension>,...]`` - File Extensions
to Include::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified file extension
or extensions in the include list.

``[-file-extensions-to-exclude <File extension>,...]`` - File Extensions
to Exclude::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified file extension
or extensions in exclude list.
```

```
``[-is-file-extension-check-on-directories-enabled {true|false}]`` - Is
File Extension Check on Directories Enabled [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified file extension
check on directories. If set to true, the command displays information
about scopes where file extension checks on directories is enabled. If set
to false, the command displays information about scopes where file
extension checks on directories is disabled.

``[-is-monitoring-of-objects-with-no-extension-enabled {true|false}]`` -
Is Monitoring of Objects with No Extension Enabled [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays scope information only
about the FPolicy policy or policies that use the specified monitoring of
objects with no extension setting. If set to true, the command displays
information about scope of policy or policies for which monitoring of
objects with no extension is enabled.

== Examples

The following example displays scope information about FPolicy policies.
```

cluster1::> vserver fpolicy policy scope show
Vserver Policy Extensions Extensions
Name Name Included Excluded
---------- ---------- ---------- ----------
Cluster cserver_pol - -
vs1.example.com p - -
vs1.example.com vs1_pol mp3 -
3 entries were displayed.

```
:leveloffset: -1


= vserver http-proxy commands

:leveloffset: +1


[[ID2506d3766637f004e4d47c7b8b2416f2]]
= vserver http-proxy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Create a HTTP Proxy configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver http-proxy create` command creates an HTTP proxy
configuration for a Vserver.

== Parameters

{ ``-vserver <vserver>`` - Vserver::
Use this parameter to specify the Vserver on which the HTTP proxy
configuration is to be created.
+
Note that `-vserver` and `-ipspace` are mutually exclusive options.

| ``[-ipspace <text>]`` - Ipspace }::
Use this parameter to specify the IPspace on which the HTTP proxy
configuration is to be created. The proxy is created on the system Vserver
of that IPspace.
+
Note that `-vserver` and `-ipspace` are mutually exclusive options.

``-server <text>`` - Proxy Server::
Use this parameter to specify the Fully Qualified Domain Name or IP
address of the proxy server.

``-port <integer>`` - Port::
Use this parameter to specify the port number on which the HTTP proxy
service is configured on the proxy server.

``-is-auth-enabled {true|false}`` - Authentication Enabled::
Use this parameter to specify if authentication is enabled to connect to
the HTTP proxy server. By default it is set to false. When set to true,
username and password will be prompted.

``[-skip-config-validation <true>]`` - Skip the config validation::
Use this parameter to skip the HTTP proxy configuration validation.
+
The proxy configuration is validated to verify that the specified server
is reachable and is providing an HTTP proxy service on the specified port.
+
The validation fails in the following cases:
+
```

```
  * The server is not reachable.
  * The specified port is invalid.
  * HTTP proxy service is not configured on the specified server and port.
  * Route or LIF does not exist.



== Examples

The following example creates an HTTP proxy configuration for the Vserver
vs0 with Fully Qualified Domain Name as input for the `-server` parameter:
```

cluster1::*> vserver http-proxy create -vserver vs0 -server example.com -port 222

```
The following example creates an HTTP proxy configuration with the IP
address as input for the `-server` parameter:
```

cluster1::*> vserver http-proxy create -vserver vs0 -server 192.168.0.0 -port 222

```
The following example creates an HTTP proxy configuration with `-skip-
config-validation` set to true:
```

cluster1::*> vserver http-proxy create -vserver vs0 -server 192.168.0.0 -port 222 -skip-config-validation

```
The following example creates an HTTP proxy configuration with the IPspace
specified and not the Vserver name:
```

cluster1::*> vserver http-proxy create -ipspace default -server 192.168.0.0 -port 222

```
The following example creates an HTTP proxy configuration with
authentication enabled:
```

cluster1::*> vserver http-proxy create -ipspace default -server 192.168.0.0 -port 222 -is-auth-enabled true

Enter the user name: test

Enter the password:

```
[[ID7bb69bc302d28d460cd7b13b457fe4b2]]
= vserver http-proxy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a HTTP Proxy configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver http-proxy delete` command deletes an HTTP proxy
configuration.

== Parameters

``-vserver <vserver>`` - Vserver::
Use this parameter to specify the Vserver for which the HTTP proxy
configuration is to be deleted.

== Examples

The following example deletes the HTTP proxy configuration for the Vserver
vs0:
```

cluster1::*> vserver http-proxy delete -vserver vs0

```
[[ID2a28c91c4379950d3b2c7b35d017f797]]
= vserver http-proxy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change a HTTP Proxy configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `vserver http-proxy modify` command modifies an HTTP proxy
configuration of a Vserver.

== Parameters

{ ``-vserver <vserver>`` - Vserver::
Use this parameter to specify the Vserver for which the HTTP proxy server
configuration is to be modified.
+
Note that `-vserver` and `-ipspace` are mutually exclusive options.

| ``[-ipspace <text>]`` - Ipspace }::
Use this parameter to specify the IPspace for which the HTTP proxy server
configuration is to be modified.
+
Note that `-vserver` and `-ipspace` are mutually exclusive options.

``[-server <text>]`` - Proxy Server::
Use this parameter to specify the Fully Qualified Domain Name or IP
address of the proxy server.

``[-port <integer>]`` - Port::
Use this parameter to specify the port number on which the HTTP proxy
service is configured on the proxy server.

``[-is-auth-enabled {true|false}]`` - Authentication Enabled::
Use this parameter to specify if authentication is enabled to connect to
the HTTP proxy server. By default it is set to false. When set to true,
username and password will be prompted.

``[-skip-config-validation <true>]`` - Skip the config validation::
Use this parameter to skip the HTTP proxy configuration validation.
+
The proxy configuration is validated to verify that the specified server
is reachable and is providing HTTP proxy service on the specified port.
+
The validation fails in the following cases:
+
  * The server is not reachable.
  * The specified port is invalid.
  * HTTP proxy service is not configured on the specified server and port.
  * Route or LIF does not exist.



== Examples
```

```
The following example modifies the HTTP proxy server of the Vserver vs0
with Fully Qualified Domain Name as input for `-server` parameter:
```

cluster1::*> vserver http-proxy modify -vserver vs0 -server example.com

```
[[IDb02562d5d9b32883a94b8cd02f1f550d]]
= vserver http-proxy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display HTTP Proxy configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver http-proxy show` command displays information about HTTP
proxy configurations. The IPspace field is populated only if a Vserver is
the system vserver of that IPspace. For all other Vservers, the IPspace
field is empty.

== 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.

{ ``[-vserver <vserver>]`` - Vserver::
Use this parameter to display information only about the HTTP proxy
configuration of the Vservers you specify.

``[-ipspace <text>]`` - Ipspace::
Use this parameter to display information only about the HTTP proxy
configuration of the IPspace you specify.

``[-server <text>]`` - Proxy Server::
```

```
Use this parameter to display information only about the HTTP proxy
configurations that match the the Fully Qualified Domain Name or IP
address you specify.

``[-port <integer>]`` - Port::
Use this parameter to display information only about the HTTP proxy
configurations that match the port you specify.

``[-is-auth-enabled {true|false}]`` - Authentication Enabled::
Use this parameter to display information only about the HTTP proxy
configurations that match the specified authentication enabled setting.

== Examples

The following example displays information about all HTTP proxy
configurations:
```

cluster1::*> vserver http-proxy show
Vserver IPSpace Server Port
--- --- --- ---
test test 10.0.0.0 3128
vs2 - 192.168.0.0 3128
2 entries were displayed.

```
The following example displays information about all HTTP proxy
configurations with `server` 10.0.0.0:
```

cluster1::*> vserver http-proxy show -server 10.0.0.0
Vserver IPSpace Server Port
--- --- --- ---
test test 10.0.0.0 3128

```
:leveloffset: -1


= vserver iscsi commands

:leveloffset: +1


[[ID6a498004b066a0253d662433883a35b2]]
= vserver iscsi create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Create a Vserver's iSCSI service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates an iSCSI target for a specified Vserver. By default
the system creates a default iSCSI target name with the status-admin set
to enabled. Until you create an iSCSI service, iSCSI initiators cannot log
into the Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver for the iSCSI service.

``[-target-name <text>]`` - Target Name [.advaced-priv]#(privilege:
advanced)#::
Specifies a iSCSI target name of a Vserver. This name is unique and is not
case sensitive. The target name must conform to this format iqn.1995-
08.com.example:string and the following rules:
+
  * Contains up to 128 bytes.
  * Contains alphanumeric characters. The period ".", hyphen "-", and
colon ":" are acceptable.
  * Does not contain the underscore character "_".



``[-target-alias <text>]`` - Target Alias::
Specifies an iSCSI target alias name of a Vserver. The maximum number of
characters for an alias name is 128. The alias default name is the Vserver
name.

``[-status-admin {down|up}]`` - Administrative Status::
Specifies the administrative status of the iSCSI service of a Vserver. If
you set this parameter to up, the command creates an iSCSI service with
the administrative status of up. If you set this parameter to down, the
command creates an iSCSI service with the administrative status of down.

``[-retain-timeout <integer>]`` - RFC3720 DefaultTime2Retain Value (in
sec) [.advaced-priv]#(privilege: advanced)#::
Specifies the wait time before an active task reassignment is possible
after an unexpected connection termination. For example, a value of 0
```

```
means that the connection or task state is immediately discarded by the
target. The default is 20 seconds.

``[-login-timeout <integer>]`` - Login Phase Duration (in sec) [.advaced-
priv]#(privilege: advanced)#::
Specifies the login phase duration. The default is 15 seconds.

``[-max-conn-per-session <integer>]`` - Max Connections per Session
[.advaced-priv]#(privilege: advanced)#::
Specifies the maximum number of connections per session that a target can
accept. The default is 4 connections.

``[-max-ios-per-session <integer>]`` - Max Commands per Session [.advaced-
priv]#(privilege: advanced)#::
Specifies the maximum number of commands per session that a target can
accept. The default is 128 commands per session.

``[-tcp-window-size <integer>]`` - TCP Receive Window Size (in bytes)
[.advaced-priv]#(privilege: advanced)#::
Specifies the TCP receive window size (in bytes). The default is 131,400
bytes.

``[-f, -force <true>]`` - Allow Non-Vendor Target Name [.advaced-
priv]#(privilege: advanced)#::
Force the command to accept a target name that would normally be rejected
as invalid.

== Examples
```

cluster1::> vserver iscsi create -vserver vs_1

```
Creates the iSCSI service for Vserver vs_1.




[[ID646f98e371aab0c6d5cb9cc1daff2d51]]
= vserver iscsi delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Vserver's iSCSI service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command deletes the iSCSI service from a Vserver.

NOTE: You must first disable the service with the command
xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify] with
"-status-admin down" before you can delete the service.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver for the iSCSI service.

== Examples
```

cluster1::> vserver iscsi delete -vserver vs_1

```
Deletes the iSCSI service for Vserver vs_1.

== Related Links

* xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify]
```

```
[[IDe97427e603e78238de031362daf1ba5a]]
= vserver iscsi modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a Vserver's iSCSI service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command modifies the configuration for an iSCSI service.

Modifications take effect immediately after you execute the command.
Making modifications to your service can result in traffic loss on a live
system. Call technical support if you are unsure of the possible
consequences.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver for the iSCSI service.

``[-target-name <text>]`` - Target Name [.advaced-priv]#(privilege:
advanced)#::
Specifies an iSCSI target name of a Vserver. This name is unique and is
not case sensitive. The target name must conform to this format iqn.1995-
08.com.example:string and the following rules:
+
  * Contains up to 128 bytes.
  * Contains alphanumeric characters. The period ".", hyphen "-", and
colon ":" are acceptable.
  * Does not contain the underscore character "_".


+
NOTE: The iSCSI service must be down in order to change the target name.
+


{ ``[-target-alias <text>]`` - Target Alias::
Specifies the new target alias of the iSCSI service.
```

```
| ``[-c, -clear <true>]`` - Clear the Target Alias }::
Clears the current target alias from the iSCSI service configuration.

``[-status-admin {down|up}]`` - Administrative Status::
Specifies the configured administrative status of a service. If you set
this parameter to up, the iSCSI service begins to accept login requests
from iSCSI initiators. If you set this parameter to down, iSCSI initiators
cannot log in.

``[-retain-timeout <integer>]`` - RFC3720 DefaultTime2Retain Value (in
sec) [.advaced-priv]#(privilege: advanced)#::
Specifies the wait time before active task reassignment is possible after
an unexpected connection termination. For example, a value of 0 means that
the connection or task state is immediately discarded by the target.

``[-login-timeout <integer>]`` - Login Phase Duration (in sec) [.advaced-
priv]#(privilege: advanced)#::
Specifies maximum time the login phase remains active until the iSCSI
target terminates the connection.

``[-max-conn-per-session <integer>]`` - Max Connections per Session
[.advaced-priv]#(privilege: advanced)#::
Specifies the maximum number of connections per session that the iSCSI
target can accept.

``[-max-ios-per-session <integer>]`` - Max Commands per Session [.advaced-
priv]#(privilege: advanced)#::
Specifies the maximum number of commands per session that the iSCSI target
can accept.

``[-tcp-window-size <integer>]`` - TCP Receive Window Size (in bytes)
[.advaced-priv]#(privilege: advanced)#::
Specifies the TCP receive window size (in bytes).
+
A change to the TCP receive window size value takes effect for all network
interfaces when you restart the iSCSI service for the Vserver as follows:
```

vserver iscsi stop -vserver <vserver name>
vserver iscsi start -vserver <vserver name>

```
If you change an individual network interface from up to down back to up,
as follows, the new value for TCP receive window size takes effect for
that network interface:
```

network interface modify -vserver <vserver name> -lif <LIF name> -status-admin down
network interface modify -vserver <vserver name> -lif <LIF name> -status-admin up

```
``[-f, -force <true>]`` - Allow Non-Vendor Target Name [.advaced-
priv]#(privilege: advanced)#::
Force the command to accept a target name that would normally be rejected
as invalid.

== Examples
```

cluster1::> vserver iscsi modify -vserver vs_1 -status-admin down

```
Modifies the status-admin of the iSCSI service for Vserver vs_1 to down.




[[IDf078f3276eca508d51680980ae295b15]]
= vserver iscsi show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a Vserver's iSCSI configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays the current configuration of the iSCSI service.

== 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.
```

```
``[-vserver <Vserver Name>]`` - Vserver::
Selects the iSCSI services for the Vserver that matches the parameter
value.

``[-target-name <text>]`` - Target Name::
Selects the iSCSI services with a target name that matches the parameter
value.

``[-target-alias <text>]`` - Target Alias::
Selects the iSCSI services with a target alias that matches the parameter
value.

``[-status-admin {down|up}]`` - Administrative Status::
Selects the iSCSI services with a configured status that matches the
parameter value.

``[-retain-timeout <integer>]`` - RFC3720 DefaultTime2Retain Value (in
sec) [.advaced-priv]#(privilege: advanced)#::
Selects the iSCSI services with a wait time that matches the parameter
value. The wait time is the amount of time before active task reassignment
is possible after an unexpected connection termination.

``[-login-timeout <integer>]`` - Login Phase Duration (in sec) [.advaced-
priv]#(privilege: advanced)#::
Selects the iSCSI services with a login phase duration that matches the
parameter value.

``[-max-conn-per-session <integer>]`` - Max Connections per Session
[.advaced-priv]#(privilege: advanced)#::
Selects the iSCSI services with a maximum connection per session that
matches the parameter value.

``[-max-ios-per-session <integer>]`` - Max Commands per Session [.advaced-
priv]#(privilege: advanced)#::
Selects the iSCSI services with a maximum number of commands per session
that matches the parameter value.

``[-tcp-window-size <integer>]`` - TCP Receive Window Size (in bytes)
[.advaced-priv]#(privilege: advanced)#::
Selects the iSCSI services with a TCP receive window size (in bytes) that
matches the parameter value.

== Examples
```

cluster1::> vserver iscsi show
Target Target Status

Vserver Name Alias Admin
---------- ---------- ---------- ------
vs_1 iqn.1992-08.com.example:sn.c7c82a22bf9f11df83e5123478563412:vs.2
vs_1_alias up
1 entries were displayed.

cluster1::> vserver iscsi show -instance
Vserver: vs_1
Target Name: iqn.1992-08.com.example:sn.c7c82a22bf9f11df83e5123478563412:vs.2

The following is the output of the show command at the advanced privilege level:
Target Alias: vs_1_alias
Administrative Status: up
1 entries were displayed.

```
Displays the output of the show command at the admin privilege level.
```

cluster1::*> vserver iscsi show
Target Target Status
Vserver Name Alias Admin
---------- ---------- ---------- ------
vs_1 iqn.1992-08.com.example:sn.c7c82a22bf9f11df83e5123478563412:vs.2
vs_1_alias up
1 entries were displayed.

```
Displays the output of the show command at the advanced privilege level.
```

cluster1::*> vserver iscsi show -instance
Vserver: vs_1
Target Name: iqn.1992-08.com.example:sn.c7c82a22bf9f11df83e5123478563412:vs.2
Target Alias: vs_1_alias
Administrative Status: up
RFC3720 DefaultTime2Retain Value (in sec): 20
Login Phase Duration (in sec): 15
Max Connections per Session: 4
Max I/O per Session: 128
TCP Window Size all Sessions (in bytes): 131400

```
Displays the detailed entries for all entries.




[[IDc28a89533fc23338445eba3923900e80]]
= vserver iscsi start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Starts the iSCSI service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command starts the iSCSI service of a Vserver. You can also use
xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify] with
"-status-admin up".

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver for the iSCSI service.

== Examples
```

cluster1::> vserver iscsi start -vserver vs_1

```
Starts the iSCSI service for Vserver vs_1.

== Related Links

* xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify]




[[IDefe5e5f566ba64a3d2d09da60abed5db]]
= vserver iscsi stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stops the iSCSI service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Stops the iSCSI service of a Vserver. This command shuts down all active
iSCSI sessions and stops any new iSCSI sessions. You can also use
xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify] with
"-status-admin down".

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver for the iSCSI service.

== Examples
```

cluster1::> vserver iscsi stop -vserver vs_1

```
Stops the iSCSI service for Vserver vs_1.

== Related Links

* xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify]
```

```
[[IDe14baa4882ac49fb7c3a6fdae45c5419]]
= vserver iscsi command show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display active iSCSI commands

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays the status of active iSCSI commands in an iSCSI
session. If you specify an iSCSI command ID, the command shows what
commands are active in a session and is useful for initiator debugging.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display a list of active iSCSI commands that match
the Vserver name that you specify.

``[-tpgroup <text>]`` - Target Portal Group::
Use this parameter to display a list of active iSCSI commands that are
within the target portal group.

``[-tsih <integer>]`` - Target Session ID::
Use this parameter to display a list of active iSCSI commands that match
the target session ID handle that you specify.

``[-command-id <integer>]`` - Command ID::
Use this parameter to display a list of active iSCSI commands that match
the command ID that you specify.
```

```
``[-initiator-name <text>]`` - Initiator Name::
Use this parameter to display a list of active iSCSI commands that match
the initiator name that you specify.

``[-initiator-alias <text>]`` - Initiator Alias::
Use this parameter to display a list of active iSCSI commands that match
the initiator alias that you specify.

``[-isid <text>]`` - Initiator Session ID::
Use this parameter to display a list of active iSCSI commands that match
the initiator session ID that you specify.

``[-command-sub-id <integer>]`` - Command Sub ID::
Use this parameter to display a list of active iSCSI commands that match
the command sub ID that you specify.

``[-command-state <iSCSI Command States>]`` - Command State::
Use this parameter to display a list of active iSCSI commands that match
the command state that you specify.

``[-command-type {Sequenced|Imm_Taskmgmt|Imm_Other}]`` - Command Type::
If you use this parameter, the command displays a list of active iSCSI
commands that contains the specified command type. The command types
indicate:
+
  * "Sequenced" -- the system processes the commands in sequence
  * "Imm_Taskmgmt" -- the system processes the commands immediately
  * "Imm_Other" -- the system processes the commands as queued



== Examples
```

cluster1::> vserver iscsi command show -instance -vserver vs_1
server: vs_1
Target Portal Group Name: tpgroup_1
Target Session ID: 2
Command ID: 20797
Initiator Name: iqn.1993-08.org.debian:01:fa752b8a5a3a
Initiator Alias: alias_1
Initiator Session ID: 00:02:3d:01:00:00
Command Sub ID: 20797
Command State: Scsicdb_Waiting_STLayer
Command Type: Sequenced

```
Displays detailed information for active iSCSI commands in Vserver vs_1.
```

```
[[ID8cbef9b9eac4b515ca2e7279dd3f2e85]]
= vserver iscsi connection show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display active iSCSI connections

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays iSCSI connection information within a session. If
you do not specify a connection, the command displays all information for
all connections.

An active iSCSI session can contain one or multiple iSCSI connections. If
an iSCSI connection has not completed the iSCSI login sequence, the iSCSI
session might not contain iSCSI connections.

This command gives real-time status of connection activity. You can use
the parameters header-digest-enabled and data-digest-enabled to
troubleshoot performance problems.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display iSCSI connections that match the Vserver
that you specify.

``[-tpgroup <text>]`` - Target Portal Group::
```

```
Use this parameter to display iSCSI connections that match the target
portal group that you specify.

``[-tsih <integer>]`` - Target Session ID::
Use this parameter to display iSCSI connections that match the target
session ID that you specify.

``[-connection-id <integer>]`` - Connection ID::
Use this parameter to display iSCSI connections that match the connection
ID that you specify.

``[-connection-state <iSCSI Connection State>]`` - Connection State::
Use this parameter to display iSCSI connections that match the connection
state you specify.

``[-has-session {true|false}]`` - Connection Has session::
Specifies if a session is established for a connection. If you enter this
command using the parameter without a value, it is set to true, and the
command displays all connections that have an established session. If you
set this parameter to false, the command displays all connections that do
not have established sessions.

``[-lif <text>]`` - Logical interface::
Use this parameter to display iSCSI connections that match the logical
interface that you specify.

``[-tpgroup-tag <integer>]`` - Target Portal Group Tag::
Use this parameter to display iSCSI connections that use the target portal
group tag that you specify.

``[-local-address <text>]`` - Local IP Address::
Use this parameter to display iSCSI connections that use the local IP
address that you specify.

``[-local-ip-port <integer>]`` - Local TCP Port::
Use this parameter to display iSCSI connections that use the local TCP
port that you specify.

``[-authentication-method {CHAP|deny|none}]`` - Authentication Type::
Use this parameter to display iSCSI connections that match the
authentication type that you specify. CHAP requires password validation.
Deny does not allow connections. None allows all connections.

``[-data-digest-enabled {true|false}]`` - Data Digest Enabled::
Specifies if data digest is enabled for a connection. If you enter this
command using the parameter without a value, it is set to true, and the
command displays all connections that support data digest. If you set this
```

```
parameter to false, the command displays all connections that do not
support data digest.

``[-header-digest-enabled {true|false}]`` - Header Digest Enabled::
Specifies if header digest is supported. If you enter this command using
the parameter without a value, it is set to true, and the command shows
all connections that support header digest. If you set this parameter to
false, the command displays all connections that do not support header
digest.

``[-rcv-window-size <integer>]`` - TCP/IP Recv Size::
Use this parameter to display iSCSI connections that match the specified
negotiated size of the TCP/IP receive window in bytes.

``[-initiator-mrdsl <integer>]`` - Initiator Max Recv Data Length::
Use this parameter to display iSCSI connections that match the maximum
length of message that the initiator can receive.

``[-remote-address <text>]`` - Remote IP address::
Use the parameter to display iSCSI connections that match the IP address
of the initiator that you specify.

``[-remote-ip-port <integer>]`` - Remote TCP Port::
Use this parameter to display iSCSI connections that match the specified
TCP port of initiator that you specify.

``[-target-mrdsl <integer>]`` - Target Max Recv Data Length::
Use this parameter to display iSCSI connections that match the maximum
message size that a target can receive.

== Examples
```

cluster1::> vserver iscsi connection show -vserver vs1
Tpgroup Conn Local Remote TCP Recv
Vserver Name TSIH ID Address Address Size
--- --- --- --- --- --- ---
vs1 vs1.iscsi 6 0 10.63.8.163 10.60.141.65 131400
vs1 vs1.iscsi 7 0 10.63.8.163 10.62.8.75 131400
2 entries were displayed.

```
Displays connection information on Vserver vs1.




[[ID54788db5ec5f3600af1a58c9963bf91]]
= vserver iscsi connection shutdown
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Shut down a connection on a node

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command shuts down a specified iSCSI connection within a session. If
you want to shut down all iSCSI connections in a session, use the
xref:{relative_path}vserver-iscsi-session-shutdown.html[vserver iscsi
session shutdown] command.

== Parameters

``-vserver <Vserver Name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Specifies the Vserver.

``-tpgroup <text>`` - Target Portal Group [.advaced-priv]#(privilege:
advanced)#::
Specifies the target portal group that contains the connection you want to
shut down.

``-tsih <integer>`` - Target Session ID [.advaced-priv]#(privilege:
advanced)#::
Specifies the target session ID that you want to shut down.

``-connection-id <integer>`` - Connection ID [.advaced-priv]#(privilege:
advanced)#::
Specifies the connection ID that you want to shut down.

== Examples
```

cluster1::*> vserver iscsi connection shutdown -vserver vs_1 -tpgroup tpgroup_1 -tsih 4 -connection-id 0

```
Forces the shutdown of an iSCSI connection with the connection ID of 0 on
Vserver vs_1 in tpgroup tpgroup_1, target session ID 4.

== Related Links

* xref:{relative_path}vserver-iscsi-session-shutdown.html[vserver iscsi
session shutdown]




[[ID79b4ca34afbe96e7a2981853ce530e08]]
= vserver iscsi initiator show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display iSCSI initiators currently connected

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays a list of active initiators currently connected to a
specified Vserver.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display the active initiators that match the Vserver
that you specify.
```

```
``[-tpgroup <text>]`` - Target Portal Group::
Use this parameter to display the active initiators that match the name of
the target portal group that you specify.

``[-tsih <integer>]`` - Target Session ID::
Use this parameter to display the active initiators that match the target
session ID you that specify.

``[-initiator-name <text>]`` - Initiator Name::
Use this parameter to display the active initiators that match the
initiator name that you specify.

``[-initiator-alias <text>]`` - Initiator Alias::
Use this parameter to display the active initiators that match the alias
name that you specify.

``[-tpgroup-tag <integer>]`` - TPGroup Tag::
Use this parameter to display the active initiators that match the target
portal group tag that you specify.

``[-isid <text>]`` - Initiator Session ID::
Use this parameter to display the active initiators that match the
initiator session ID that you specify.

``[-igroup <text>,...]`` - Igroup Name::
Use this parameter to display the active initiators that match the
initiator group that you specify.

== Examples
```

cluster1::> vserver iscsi initiator show -vserver vs_1
Tpgroup Initiator
Vserver Name TSIH Name ISID IGroup
------- -------- ---- ---------- ---------- ----------
vs_1 vs_1.iscsi 6 iqn.1994-05.com.redhat:6ed6dfb0489e
00:02:3d:03:00:00 -
vs_1 vs_1.iscsi 7 iqn.1993-08.org.debian:01:fa752b8a5a3a
00:02:3d:01:00:00 igroup_1
2 entries were displayed.

```
Displays the active initiator information on Vserver vs_1.




[[ID0979ca76423e73792809a9a7c570a3e1]]
= vserver iscsi interface disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable the specified interfaces for iSCSI service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command disables the specified logical interfaces for an iSCSI
service. Once disabled, all subsequent attempts to establish new iSCSI
connections over the logical interface will fail.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

{ ``-lif <lif-name>,...`` - Logical Interface::
Specifies the logical interfaces on a Vserver you want to disable.

| ``-a, -all <true>`` - All }::
Specifies that all logical interfaces on the Vserver are disabled.

``[-f, -force <true>]`` - Force::
When set to true, forces the termination of any active iSCSI sessions
without prompting you for a confirmation.

== Examples
```

cluster1::> vserver iscsi interface disable -vserver vs_1 -lif vs_1.iscsi

```
Disables the iscsi logical interface vs_1.iscsi on Vserver vs_1.




[[ID80101f2e862dd38a02e2ed236efcadab]]
= vserver iscsi interface enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable the specified interfaces for iSCSI service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command enables specified logical interfaces for iSCSI Vserver
service. Once enabled, your system accepts new iSCSI connections and
services iSCSI requests over the newly enabled logical interfaces.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

{ ``-lif <lif-name>,...`` - Logical Interface::
Specifies the logical interfaces on a Vserver that you want to enable.

| ``-a, -all <true>`` - All }::
When set to true, all logical interfaces are enabled. If you use this
parameter without a value, it is set to true, and the command enables all
logical interfaces.

== Examples
```

cluster1::> vserver iscsi interface enable -vserver vs_1 -lif vs_1.iscsi

```
Enables the iscsi logical interface vs_1.iscsi on Vserver vs_1.




[[IDad46434a0849443460a9674fa362f41d]]
= vserver iscsi interface modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify network interfaces used for iSCSI connectivity

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver iscsi interface modify` command modifies the iSCSI specific
configuration for an iSCSI LIF.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

``-lif <lif-name>`` - Logical Interface::
Use this parameter to specify the logical interface on a Vserver that you
want to modify.

``[-sendtargets-fqdn <text>]`` - iSCSI Discovery SendTargets FQDN
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the Fully Qualified Domain Name (FQDN) to
return during an iSCSI Discovery SendTargets operation. To clear the FQDN,
set this parameter to "". If unset, the IP address of the LIF is used in
iSCSI SendTargets discovery. + This is not part of iSNS and will not
affect the iSNS configuration.

== Examples

The following example modifies the sendtargets-fqdn of the iSCSI LIF
vs1_iscsi1 for Vserver vs1 to myhost.example.com.
```

cluster1::> vserver iscsi interface modify -vserver vs1 -lif vs1_iscsi1 -sendtargets_fqdn myhost.example.com

```
[[ID6376b6400787cff27fd5ddaf01fa6531]]
= vserver iscsi interface show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show network interfaces used for iSCSI connectivity

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command shows the iSCSI logical interfaces for a specified Vserver.
If you do not specify any of the parameters, the command displays all of
the interfaces on a Vserver.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display iSCSI logical interfaces that match the
Vserver that you specify.

``[-lif <lif-name>]`` - Logical Interface::
Use this parameter to display iSCSI logical interfaces that that you
specify.

``[-status-admin {up|down}]`` - Administrative Status::
Specifies the configured status of the logical interface. If you set this
parameter to up, the command displays all iSCSI logical interfaces with
the administrative status of up. If you set this parameter to down, the
command displays all the iSCSI logical interfaces with the administrative
status of down.
```

```
``[-status-oper {up|down}]`` - Operational Status::
Specifies the current status of the logical interface. If you set this
parameter to up, the command displays all the iSCSI logical interfaces
with the operational status of up. If you set this parameter to down, the
command displays all the iSCSI logical interfaces with the operational
status of down.

``[-enabled {true|false}]`` - Enabled::
Specifies if this logical unit is enabled for iSCSI service. If you enter
this command without a parameter, its effective value is true, and the
command displays all the enabled iSCSI logical interfaces.

``[-address <IP Address>]`` - IP Address::
Use this parameter to display iSCSI logical interfaces that match the IP
address that you specify.

``[-ip-port <integer>]`` - IP Port Number::
Use this parameter to display iSCSI logical interfaces that match IP port
number for the logical interface that you specify.

``[-curr-node <nodename>]`` - Current Node::
Use this parameter to display iSCSI logical interfaces that match current
node that you specify.

``[-curr-port {<netport>|<ifgrp>}]`` - Current Port::
Use this parameter to display iSCSI logical interfaces that match
specified current physical port that you specify.

``[-is-home {true|false}]`` - Is Home::
Specifies if the node hosting the logical interface is the initially
configured node. If you use this command without using this parameter, it
is set to true, and the command displays all iSCSI interfaces that are on
the initially configured node.

``[-tpgroup <text>]`` - TPGroup Name::
Use this parameter to display iSCSI logical interfaces that match the
target portal group name that you specify.

``[-t, -tpgroup-tag <integer>]`` - TPGroup Tag::
Use this parameter to display iSCSI logical interfaces that match the
target portal group tag that you specify.

``[-relative-port-id <integer>]`` - Relative Port ID::
Use this parameter to display the iSCSI logical interface that matches the
relative target port ID that you specify. The system assigns each logical
```

```
interfaces and target portal group a relative target port ID that is
Vserver unique. You cannot change this ID.

``[-sendtargets-fqdn <text>]`` - iSCSI Discovery SendTargets FQDN
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display the iSCSI logical interfaces that match the
iSCSI Discovery SendTargets Fully Qualified Domain Name (FQDN) that you
specify.

``[-home-node <nodename>]`` - Home Node::
Use this parameter to display iSCSI logical interfaces that match home
node that you specify.

``[-home-port {<netport>|<ifgrp>}]`` - Home Port::
Use this parameter to display iSCSI logical interfaces that match
specified home physical port that you specify.

== Examples

The following example displays information for logical interfaces on
Vserver vs_1.
```

cluster1::> vserver iscsi interface show -vserver vs_1
Logical Status IP Curr Curr
Vserver Interface TPGT Admin/Oper Address Node Port Enabled
---------- ---------- ---- ---------- --------------- ----------- ---- -------
vs_1 vs_1.iscsi 1027 up/up 10.63.8.165 node1 e0c true
vs_1 vs_1.iscsi2
1028 up/up 10.63.8.166 node1 e0c true
2 entries were displayed.

```
The following example displays the logical interface vs_1.iscsi with the
relative target port ID of 1.
```

cluster1::> vserver iscsi interface show -vserver vs_1 -relative-port-id 1
Logical Status IP Curr Curr
Vserver Interface TPGT Admin/Oper Address Node Port Enabled
---------- ---------- ---- ---------- --------------- ----------- ---- -------
vs_1 vs_1.iscsi 1027 up/up 10.63.8.165 node1 e0c true

```
[[ID8bd1ab933b74ac5483a8531a8d60a99b]]
= vserver iscsi interface accesslist add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Add the iSCSI LIFs to the accesslist of the specified initiator

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command adds network interfaces to an access list for a specified
initiator. An access list ensures that an initiator only logs in with IP
addresses associated with the interfaces defined in the access list.

You can restrict an initiator to certain network interfaces to improve
performance and security. Access lists are useful where a particular
initiator cannot access all of the network interfaces on a node.

Access list policies are based on the interface name. The accesslist rules
are:

* If you disable the network interface for iSCSI through the
xref:{relative_path}vserver-iscsi-interface-disable.html[vserver iscsi
interface disable] command, for example, the network interface is not
accessible to any initiator regardless of any access lists in effect.
* If an initiator does not have an access list, that initiator can access
any iSCSI-enabled network interface.
* If an initiator has an access list, that initiator can only login to
network interfaces in its access list. Additionally, the initiator cannot
discover any IP addresses that are not on this access list. If an
initiator sends an iSCSI sendtargets request, the node responds with a
list of IP addresses for iSCSI data logical interfaces that are in its
access list.
* If an initiator does not have an access list, you automatically create
an access list when you issue the `vserver iscsi interface accesslist add`
command.
* If you remove all the interfaces from the access list of an initiator
with the xref:{relative_path}vserver-iscsi-interface-accesslist-
remove.html[vserver iscsi interface accesslist remove] command, the
accesslist is also deleted.
* Creating or modifying access list requires that initiator log out and
log back in before changes take effect.



When you use the add or remove commands, the system warns you if an iSCSI
session could be affected.
```

```
NOTE: You will not affect any iSCSI sessions if you use the -a parameter
when adding or removing all interfaces.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver name.

``-initiator-name <text>`` - Initiator Name::
Specifies the initiator you want to add to the access list.

{ ``-lif <lif-name>,...`` - Logical Interface::
Specifies the lif you want to add to an access list.

| ``-a, -all <true>`` - All }::
If you use this parameter without a value, it is set to true, and the
command adds all iSCSI data logical interfaces for a vserver to an
initiator's accesslist. If the initiator does not have an accesslist, the
system creates a new accesslist.

``[-f, -force <true>]`` - Force::
If you use this parameter without a value, it is set to true, and the
command does not prompt you when an active iSCSI service or any active
iSCSI data logical interfaces could be affected. If you do not use this
parameter, the command prompts for confirmation if the iSCSI service is
active or if any active data logical interfaces would be affected.

== Examples
```

cluster1::> vserver iscsi interface accesslist add -vserver vs_1 -initiator-name iqn.1992-08.com.example:abcdefg -a

```
Adds the initiator iqn.1992-08.com.example:abcdefg on Vserver vs_1 for all
iSCSI data logical interfaces in vs_1.

== Related Links

* xref:{relative_path}vserver-iscsi-interface-disable.html[vserver iscsi
interface disable]
* xref:{relative_path}vserver-iscsi-interface-accesslist-
remove.html[vserver iscsi interface accesslist remove]
```

```
[[ID2078b82daa53f6b38d6a4829a472a36f]]
= vserver iscsi interface accesslist remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove the iSCSI LIFs from the accesslist of the specified initiator

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command removes network interfaces from an access list for a
specified initiator. The system removes the access list when the list is
empty. When you remove a network interface from an initiator, this action
could result in the shutdown of active sessions.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver name.

``-initiator-name <text>`` - Initiator Name::
Specifies the initiator that you want to remove logical interfaces from.

{ ``-lif <lif-name>,...`` - Logical Interface::
Specifies the logical interface you want to remove.

| ``-a, -all <true>`` - All }::
If you use this parameter without a value, it is set to true, and the
command removes all of the iSCSI data logical interfaces from an
initiator's accesslist. If you remove all the network interfaces from an
access list, the system removes the access list.

``[-f, -force <true>]`` - Force::
If you use this parameter without a value, it is set to true, and the
command does not prompt you when an active iSCSI service or any active
iSCSI data logical interfaces could be affected. If you do not use this
parameter, the command prompts for confirmation if the iSCSI service is
active or if any active data logical interfaces would be affected.

== Examples
```

cluster1::> vserver iscsi interface accesslist remove -vserver vs_1 -initiator-name iqn.1992-08.com.example:abcdefg -a

```
Removes all the network interfaces from the access list for initiator
iqn.1992-08.com.example:abcdefg on Vserver vs_1.




[[ID1116857cd62c218c31c3b1a06642e1c6]]
= vserver iscsi interface accesslist show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show accesslist of the initiators for iSCSI connectivity

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays an access list for an initiator. An access list is a
list of logical interfaces that an initiator can use for iSCSI logins. The
system records the access lists as part of the node configuration and
preserves the access lists during reboots.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display the access lists that match the Vserver name
that you specify.

``[-initiator-name <text>]`` - Initiator Name::
Use this parameter to display the access lists that match the initiator
```

```
that you specify.

``[-lif <lif-name>]`` - Logical Interface::
Use this parameter to display the access lists that match the logical
interface that you specify.

== Examples
```

cluster1::> vserver iscsi interface accesslist show -vserver vs1
Vserver Initiator Name Logical Interface
--- --- ---
vs1 iqn.2010-01.com.example:aaaaa isw1
isw2
iqn.2010-01.com.example:aaabb isw1
isw2
4 entries were displayed.

```
Displays the access lists for vserver vs1.




[[IDab1a0e80add2ae7e4b3b7bfd0b6633f7]]
= vserver iscsi isns create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Configure the iSNS service for the Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command creates and starts an iSNS service with the IP address of the
iSNS server.

NOTE: A Vserver management LIF must exist before you can create an iSNS
service. This LIF is used to communicate with the iSNS server. To create a
Vserver management LIF, use the xref:{relative_path}network-interface-
create.html[network interface create] command, with `-role```_data_`` and
`-data-protocol```_none_`` .
```

```
== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the iSNS service that you want to create.

``-address <IP Address>`` - iSNS Server IP Address::
Specifies the IP address of the iSNS server. Both IPv4 and IPv6 address
families are supported. The address family must be the same as that of the
vserver management LIF.
+
NOTE: A default route must exist for the specified vserver. To create a
route, use the `network routing-groups route create` command. To view
existing routes, use the `network routing-groups route show` command.
+


``[-status-admin {down|up}]`` - Administrative Status::
Specifies the administrative status of the iSNS service of a Vserver. If
you set this parameter to up, the iSNS service starts for the Vserver and
registers with the configured iSNS server. If you set this parameter to
down, the Vserver loses its ability to register with the iSNS server and
to be discovered by iSNS clients.

``[-force <true>]`` - Force::
`vserver iscsi isns create` fails if vserver management LIF is not
configured. When you set this option to "true," you create an iSNS service
on a Vserver even if the vserver does not have a vserver management LIF.

== Examples
```

cluster1::> vserver iscsi isns create -vserver vs_1 -address 10.60.1.1 -status-admin up

```
Creates the iSNS service for Vserver vs_1 using the IPv4 address.
```

cluster1::> vserver iscsi isns create -vserver vs_1 -address fd20:8b1e:b255:840b:a0df:565b:19b5:4d06 -status -admin up

```
Creates the iSNS service for Vserver vs_1 using the IPv6 address.

== Related Links

* xref:{relative_path}network-interface-create.html[network interface
create]




[[ID67e4be3e8072cc2794a73b8a795f99e6]]
= vserver iscsi isns delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove the iSNS service for the Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command deletes the iSNS service for the Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the iSNS service that you want to delete.

== Examples
```

cluster1::> vserver iscsi isns delete -vserver vs_1

```
Deletes the iSNS service for Vserver vs_1.




[[IDd4fe524be3f9e662f5b3790dca3244fa]]
= vserver iscsi isns modify
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the iSNS service for the Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command modifies the configuration of an iSNS service.

Modifications take effect immediately after you execute the command.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the iSNS service that you want to modify.

``[-address <IP Address>]`` - iSNS Server IP Address::
Specifies the IP address of the iSNS server. Both IPv4 and IPv6 address
families are supported. The address family must be the same as that of the
vserver management LIF.
+
NOTE: A default route must exist for the specified vserver. To create a
route, use the `network routing-groups route create` command. To view
existing routes, use the `network routing-groups route show` command.
+


``[-status-admin {down|up}]`` - Administrative Status::
Specifies the administrative status of the iSNS service of a Vserver. If
you set this parameter to up, the iSNS service starts for the Vserver, and
registers with the configured iSNS server. If you set this parameter to
down, the Vserver loses its ability to register with the iSNS server and
to be discovered by iSNS clients.

``[-force <true>]`` - Force::
`vserver iscsi isns modify` fails to modify the iSNS server address if
vserver management LIF is not configured. When you set this option to
"true," you can modify the iSNS service on a Vserver even if the vserver
does not have a vserver management LIF.

== Examples
```

cluster1::> iscsi isns modify -vserver vs_1 -status-admin up

```
Modifies the status-admin of the iSNS service for Vserver vs_1 to up.




[[ID5c46daa2610a779822e3d1f758c0a45a]]
= vserver iscsi isns show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show iSNS service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Shows the iSNS service configuration.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the iSNS services that match the Vserver
name that you specify.

``[-address <IP Address>]`` - iSNS Server IP Address::
Use this parameter to display the iSNS services that match the IP address
of the iSNS server that you specify.

``[-status-admin {down|up}]`` - Administrative Status::
Use this parameter to display the iSNS services that match the configured
status of the service that you specify.
```

```
``[-entity-id <text>]`` - iSNS Server Entity Id::
Use this parameter to display the iSNS services that match the configured
iSNS server entity-id that you specify.

``[-last-successful-update <MM/DD/YYYY HH:MM:SS>]`` - Last Successful
Update::
Use this parameter to display the iSNS services that match the time of the
last successful attempt.

``[-last-update-attempt <MM/DD/YYYY HH:MM:SS>]`` - Last Update Attempt::
Use this parameter to display the iSNS services that match the time of the
last update attempt.

``[-last-update-result <isnsErrors>]`` - Last Update Result::
Use this parameter to display the iSNS services that match the result of
the last update attempt.

== Examples
```

cluster1::> vserver iscsi isns show
Vserver iSNS Server Entity Identifier iSNS Server IP Address iSNS Status
---------- ---------- ---------- ----------
iscsi_vs isns:00000044 10.229.136.188 up

```
Displays the output of the show command for all Vservers in a cluster.
```

cluster1::> vserver iscsi isns show -instance
Vserver Name: vs1
iSNS Server IP Address: 10.72.19.11
Administrative Status: up
iSNS Server Entity Id: isns.0000001c
Last Successful Update: 11/12/2011 10:18:45
Last Update Attempt: 11/12/2011 10:18:45
Last Update Result: iSNS_Ok

Vserver Name: vs2
iSNS Server IP Address: 10.72.16.13
Administrative Status: up
iSNS Server Entity Id: isns.0000001b
Last Successful Update: 11/12/2011 13:38:05
Last Update Attempt: 11/12/2011 13:38:05
Last Update Result: iSNS_Ok

2 entries were displayed.

```
Displays the details for all Vservers in a cluster.




[[ID7a290e6f18c9fd4f2cee7ee041de76d0]]
= vserver iscsi isns start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Starts the iSNS service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Starts the iSNS service. Once you start the iSNS service, the Vserver
automatically register with the iSNS server.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the iSNS service that you want to start.

== Examples
```

cluster1::> vserver iscsi isns start -vserver vs_1

```
Starts the iSNS service for Vserver vs_1.




[[ID74c0c9a7e01cda20179595db5f5c268e]]
= vserver iscsi isns stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stops the iSNS service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Stops the iSNS service. Once you stop the iSNS service, the Vserver loses
the ability to register with the iSNS server and to be discovered by iSNS
clients.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the iSNS service that you want to stop.

== Examples
```

cluster1::> vserver iscsi isns stop -vserver vs_1

```
Stops the iSNS service for Vserver vs_1.




[[ID8ea884efb09392461a0c91a26a835e10]]
= vserver iscsi isns update
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Force update of registered iSNS information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Forces an update of the registration information with the iSNS server.
Normally, the system checks for iSNS configuration changes on the Vserver
every few minutes and automatically sends updates to the iSNS server.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the iSNS service that you want to update.

== Examples
```

cluster1::> vserver iscsi isns update -vserver vs_1

```
Updates the iSNS server registration for Vserver vs_1.




[[ID82df8186d760e014cdb9e807831df7e]]
= vserver iscsi security add-initator-address-ranges
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add IP Address Ranges

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Add IP address ranges to an existing iSCSI security entry

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

``-i, -initiator-name <text>`` - Initiator Name::
Specifies the initiator.

``-initiator-address-ranges {<ipaddr>|<ipaddr>-<ipaddr>}`` - Initiator IP
Address Ranges::
Specifies one or more initiator source IP address range. The IPv4 or IPv6
address range contains a start address and an end address. The start and
end addresses themselves are included in the range.
+
An example of a valid IPv4 address range is: '192.168.1.100-
192.168.1.150'.
+
An example of a valid IPv6 address range is: '2001:db8::1000:1-
2001:db8::1000:50'.

== Examples
```

cluster1::> vserver iscsi security add-initator-address-range

-vserver vs1 -initiator-name iqn.1993-08.com.example:01:e3f87c7cf2e4
-initiator-address-range 192.168.2.1-192.168.2.255

```
Adds the IP address range 192.168.2.1-192.168.2.255 to initiator iqn.1993-
08.com.example:01:e3f87c7cf2e4 for vserver vs1.




[[ID277b5cacd8e7ac4f541d53cce3538fa4]]
= vserver iscsi security create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an iSCSI authentication configuration for an initiator

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command configures the security method for an iSCSI initiator on a
Vserver. The outbound CHAP password and user name are optional. If you
want mutual authentication, you need to configure both inbound and
outbound CHAP passwords and user names.

You cannot use the same password for inbound and outbound settings.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

``-i, -initiator-name <text>`` - Initiator Name::
Specifies the initiator that you want to create a security method for. You
can use either an iqn such as iqn.1995-08.com.example:string or eui such
as eui.0123456789abcdef for the initiator.

``-s, -auth-type {CHAP|deny|none}`` - Authentication Type::
Specifies the authentication type:
+
  * CHAP - Authenticates using a CHAP user name and password.
  * none - The initiator can access the Vserver without authentication.
```

```
  * deny - The initiator cannot access the Vserver.



``[-n, -user-name <text>]`` - Inbound CHAP User Name::
Specifies the inbound CHAP user name. CHAP user names can be one to 128
bytes. A null user name is not allowed. If provided, you will be prompted
to provide the corresponding inbound CHAP password.

``[-m, -outbound-user-name <text>]`` - Outbound CHAP User Name::
Specifies the outbound CHAP user name. CHAP user names can be one to 128
bytes. If provided, you will be prompted to enter the corresponding
outbound CHAP password.

``[-initiator-address-ranges {<ipaddr>|<ipaddr>-<ipaddr>}]`` - Initiator
IP Address Ranges::
Specifies one or more initiator source IP address ranges. If this list is
empty, the initiator is allowed to log in from any IP address. The IPv4 or
IPv6 address range contains a start address and an end address. The start
and end addresses themselves are included in the range.
+
An example of a valid IPv4 address range is: '192.168.1.100-
192.168.1.150'.
+
An example of a valid IPv6 address range is: '2001:db8::1000:1-
2001:db8::1000:50'.

== Examples
```

cluster1::> vserver iscsi security create -vserver vs_1
-initiator-name eui.0123456789abcdef -auth-type CHAP -user-name bob -outbound-user-name bob2

Password: {enter password}

Outbound Password: {enter password}

```
Creates authentication method chap for initiator eui.0123456789abcdef with
inbound and outbound usernames and passwords.
```

cluster1::> vserver iscsi security create -vserver vs_1
-initiator-name iqn.1995-08.com.example:e3f87c7cf2e4 -auth-type none
-initiator-address-ranges 192.168.1.1-192.168.1.255

```
Creates authentication method for initiator iqn.1993-
08.com.example:01:e3f87c7cf2e4 with IP address validation only.
```

```
[[ID3fb8db626f922583cae71267d7c0f09f]]
= vserver iscsi security default
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Configure the default authentication settings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command defines a default iSCSI authentication method for your
Vserver. If you do not configure the initiator to use a user-defined
authentication method, the system assigns the default authentication
method automatically to the initiator. Use the
xref:{relative_path}vserver-iscsi-security-create.html[vserver iscsi
security create] command if you want to configure a user-defined
authentication method.

The outbound CHAP user name and password are optional. If you want a bi-
directional handshake, provide the outbound user name and you will be
prompted for the corresponding password.

You cannot use the same password for inbound and outbound settings.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

``-s, -auth-type {CHAP|deny|none}`` - Authentication Method::
Specifies the authentication type:
+
  * CHAP - Authenticates using a CHAP user name and password.
  * none - The initiator can access the Vserver without authentication.
  * deny - The initiator cannot access the Vserver.



``[-n, -user-name <text>]`` - Inbound CHAP User Name::
```

```
Specifies the inbound CHAP user name. CHAP user names can be one to 128
bytes. A null user name is not allowed. If provided, you will be prompted
to provide the corresponding inbound CHAP password.

{ ``[-m, -outbound-user-name <text>]`` - Outbound CHAP User Name::
Specifies the outbound CHAP user name. CHAP user names can be one to 128
bytes. If provided, you will be prompted to enter the corresponding
outbound CHAP password.

| ``[-clear-outbound <true>]`` - Clear Outbound CHAP Parameters }::
Removes the outbound user name and the outbound password information from
the default authentication method. After you clear the outbound
information, you no longer have a bi-directional handshake.

== Examples
```

cluster1::> vserver iscsi security default -vserver vs1 -auth-type chap -user-name bob -outbound-user-name bob_out

Password:

Outbound Password:

```
Sets the default authentication method to CHAP with inbound and outbound
user names and passwords.

== Related Links

* xref:{relative_path}vserver-iscsi-security-create.html[vserver iscsi
security create]




[[ID63accb37885bf93a0c36a7370125189d]]
= vserver iscsi security delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete the iSCSI authentication configuration for an initiator

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command removes the security settings for this initiator. The default
authentication setting now applies to this initiator.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

``-i, -initiator-name <text>`` - Initiator Name::
Specifies the initiator that you want to remove the authentication setting
from.

== Examples
```

cluster1::> vserver iscsi security delete -vserver vs1 -initiator-name iqn.1992-08.com.example:abcdefg

```
Deletes initiator iqn.1992-08.com.example:abcdefg on Vserver vs1 from the
authentication setting. The default authentication now applies to this
initiator.
```

```
[[ID9911bc5d11d055c4f979307a691c9dda]]
= vserver iscsi security modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the iSCSI authentication configuration for an initiator

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The command modifies an existing authentication method for an initiator.
To delete the authentication setting for an initiator, use the
xref:{relative_path}vserver-iscsi-security-delete.html[vserver iscsi
security delete] command.

The outbound CHAP password and user name are optional. If you want a bi-
directional handshake, you need to configure both inbound and outbound
CHAP passwords and user names.

You do not need to know the inbound or outbound passwords to change them.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

``-i, -initiator-name <text>`` - Initiator Name::
Specifies the initiator name that you want to modify the existing
authentication method.

``[-s, -auth-type {CHAP|deny|none}]`` - Authentication Type::
Specifies the authentication type:
+
  * CHAP - Authenticates using a CHAP user name and password.
  * none - The initiator can access the Vserver without authentication.
  * deny - The initiator cannot access the Vserver.
```

```
``[-n, -user-name <text>]`` - Inbound CHAP User Name::
Specifies the inbound CHAP user name. CHAP user names can be one to 128
bytes. A null user name is not allowed. If provided, you will be prompted
to provide the corresponding inbound CHAP password.

{ ``[-m, -outbound-user-name <text>]`` - Outbound CHAP User Name::
Specifies the outbound CHAP user name. CHAP user names can be one to 128
bytes. If provided, you will be prompted to enter the corresponding
outbound CHAP password.

| ``[-clear-outbound <true>]`` - Clear Outbound CHAP Parameters }::
Removes the outbound user name and the outbound password information from
the authentication method. After you clear the outbound information, you
no longer have a bi-directional handshake.

== Examples
```

cluster1::> vserver iscsi security modify -vserver vs_1 -initiator-name iqn.1992-08.com.example:abcdefg -auth -type chap -user-name bob -outbound-user-name bob_out

Password:

Outbound Password:

```
Changes user names and passwords for initiator iqn.1992-
08.com.example:abcdefg on Vserver vs_1.

== Related Links

* xref:{relative_path}vserver-iscsi-security-delete.html[vserver iscsi
security delete]




[[ID8f915e6c310b879f2d3e97a82bb1dc44]]
= vserver iscsi security prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Prepares the system for downgrade

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command prepares the cluster for a downgrade to an earlier version of
ONTAP. Before using this command verify that all security entries do not
have any initiator address ranges defined. This may be done by running the
command xref:{relative_path}vserver-iscsi-security-show.html[vserver iscsi
security show] ``_address-ranges_``

== Examples
```

cluster1::> vserver iscsi security prepare-to-downgrade

```
The above example will verify that the cluster is able to downgrade to a
prior release of ONTAP.

== Related Links

* xref:{relative_path}vserver-iscsi-security-show.html[vserver iscsi
security show]
```

```
[[IDb2f94187d2231c90b82a42e3ed66d093]]
= vserver iscsi security remove-initator-address-ranges
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove an IP Address Range

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Remove IP address ranges to an existing iSCSI security entry

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Specifies the Vserver.

``-i, -initiator-name <text>`` - Initiator Name::
Specifies the initiator.

``-initiator-address-ranges {<ipaddr>|<ipaddr>-<ipaddr>}`` - Initiator IP
Address Ranges::
Specifies one or more initiator source IP address range. The IPv4 or IPv6
address range contains a start address and an end address. The start and
end addresses themselves are included in the range.
+
An example of a valid IPv4 address range is: '192.168.1.100-
192.168.1.150'.
+
An example of a valid IPv6 address range is: '2001:db8::1000:1-
2001:db8::1000:50'.

== Examples
```

netapp-clus-1::> vserver iscsi security remove-initator-address-range
-vserver vs1 -initiator-name iqn.1993-08.com.example:01:e3f87c7cf2e4
-initiator-address-range 192.168.2.1-192.168.2.255

```
Removes the IP address range 192.168.2.1-192.168.2.255 to the initiator
```

```
iqn.1993-08.com.example:01:e3f87c7cf2e4 for vserver vs1.




[[IDa2b113e457318c9fe90eebfce218462c]]
= vserver iscsi security show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the current iSCSI authentication configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays the default authentication and all initiator-
specific authentication information. ONTAP authentication overrides all
other service authentication methods.

== 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.

| ``[-address-masks  ]``::
Display the list of IP Address ranges in CIDR notation that each initiator
is allowed to originate from. If this list is empty, the initiator is
allowed to log in from any IP address. The IPv4 or IPv6 address range
contains a start address and an end address. The start and end addresses
themselves are included in the range.

| ``[-address-ranges  ]``::
Display the list of IP Address ranges that each initiator is allowed to
originate from. If this list is empty, the initiator is allowed to log in
from any IP address. The IPv4 or IPv6 address range contains a start
address and an end address. The start and end addresses themselves are
included in the range.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
```

```
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display authentication information that matches the
Vserver name that you specify.

``[-i, -initiator-name <text>]`` - Initiator Name::
Use this parameter to display authentication information that matches the
initiator that you specify.

``[-s, -auth-type {CHAP|deny|none}]`` - Authentication Type::
Use this parameter to display authentication information that matches the
authentication type that you specify.

``[-n, -user-name <text>]`` - Inbound CHAP User Name::
Use this parameter to display authentication information that matches the
inbound CHAP user name that you specify.

``[-m, -outbound-user-name <text>]`` - Outbound CHAP User Name::
Use this parameter to display authentication information that matches the
outbound CHAP user name that you specify.

``[-auth-chap-policy <local>]`` - Authentication CHAP Policy::
Use this parameter to display authentication information that matches the
authentication CHAP policy that you specify.

``[-initiator-address-ranges {<ipaddr>|<ipaddr>-<ipaddr>}]`` - Initiator
IP Address Ranges::
Use this parameter to display authentication information that matches the
initiator address range that you specify. If this list is empty, the
initiator is allowed to log in from any IP address. The IPv4 or IPv6
address range contains a start address and an end address. The start and
end addresses themselves are included in the range.
+
An example of a valid IPv4 address range is: '192.168.1.100-
192.168.1.150'.
+
An example of a valid IPv6 address range is: '2001:db8::1000:1-
2001:db8::1000:50'.

``[-initiator-address-masks <IP Address/Mask>,...]`` - Initiator IP
Address Masks::
Use this parameter to display authentication information that matches the
initiator address masks that you specify. If this list is empty, the
initiator is allowed to log in from any IP address. The IPv4 or IPv6
address range contains a start address and an end address. The start and
```

```
end addresses themselves are included in the range.
+
An example of a valid IPv4 address range in CIDR notation is:
192.168.1.3/32.
+
An example of a valid IPv6 address range in CIDR notation is:
2001:db8::1000:1/128.

== Examples
```

cluster1::> vserver iscsi security show -vserver vs1
Auth Auth CHAP Inbound CHAP Outbound CHAP
Vserver Initiator Name Type Policy User Name User Name
---------- ---------- ------ --------- ---------- ----------
vs1 default none - - -
iqn.2010-12.com.example:abcdefg
CHAP local bob bob2
2 entries were displayed.

```
Displays the authentication information for Vserver vs1.
```

cluster1::> vserver iscsi security show -address-ranges -vserver vs1

Vserver Initiator Name Initiator Address Ranges
---------- ---------- ----------
vs1 iqn.2010-12.com.example:abcdefg
iqn.2010-12.com.example:hijklmn
192.168.1.100-192.168.1.150
2001:db8::1000:1-2001:db8::1000:50

2 entries were displayed.

```
Displays the initiator and their valid address ranges for Vserver vs1.
```

cluster1::> vserver iscsi security show -address-masks -vserver vs1

Vserver Initiator Name Initiator Address Ranges
---------- ---------- ----------
vs1 iqn.2010-12.com.example:abcdefg
iqn.2010-12.com.example:hijklmn
192.168.1.100/30
192.168.1.104/29
192.168.1.112/28
192.168.1.128/28
192.168.1.144/30
192.168.1.148/31
192.168.1.150/32
2001:db8::1000:1/128

2001:db8::1000:2/127
2001:db8::1000:4/126
2001:db8::1000:8/125
2001:db8::1000:10/124
2001:db8::1000:20/123
2001:db8::1000:40/124
2001:db8::1000:50/128
2 entries were displayed.

```
Displays the initiator and their valid address ranges for Vserver vs1.




[[ID91c0dbd7a76eb3ad43ed4ccf49ba3e92]]
= vserver iscsi session show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display iSCSI sessions

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays iSCSI session information. If you do not specify the
target session ID (TSIH), the command displays all session information for
the specified Vserver. If a Vserver is not specified, the command displays
all session information in the cluster. Use the
xref:{relative_path}vserver-iscsi-connection-show.html[vserver iscsi
connection show] command to display connection information. Use the
xref:{relative_path}vserver-iscsi-session-parameter-show.html[vserver
iscsi session parameter show] command to show the parameters used when
creating the session.

You can use session information for troubleshooting performance problems.

An iSCSI session can have one or multiple connections. Typically a session
has at least one connection.

Most of the parameters are read-only. However, some parameters can be
modified with the xref:{relative_path}vserver-iscsi-modify.html[vserver
iscsi modify] command.
```

```
== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display iSCSI session information that matches the
Vserver name that you specify.

``[-tpgroup <text>]`` - Target Portal Group::
Use this parameter to display iSCSI session information that matches the
target portal group name that you specify.

``[-tsih <integer>]`` - Target Session ID::
Use this parameter to display iSCSI session information that matches the
target session ID that you specify.

``[-max-ios-per-session <integer>]`` - Max Commands per Session::
Use this parameter to display iSCSI session information that matches the
maximum commands per session count you specify.

``[-data-pdu-in-order {true|false}]`` - Data PDU in Order::
Specifies if the data PDUs are in sequence order. If you enter this
command without using this parameter, it is set to true, and the command
displays all session information that supports PDUs in order. If you
provide a false value, the command displays all session information that
does not support PDUs in order.

``[-data-sequence-in-order {true|false}]`` - Data Sequence in Order::
Specifies if the data is in sequence order. If you enter this command
without using this parameter, it is set to true, and the command displays
all session information where data sequence is supported. If you provide a
false value, the command displays all session information that does not
support data sequence.

``[-default-time-to-retain <integer>]`` - Default Time to Retain::
Use this parameter to display session information that matches the retain
time that you specify. This value specifies the amount of time before
active reassignment is possible after an unexpected connection termination
```

```
or a connection reset. A value of 0 means the connection task state is
immediately discarded by the target.

``[-default-time-to-wait <integer>]`` - Default Time to Wait::
Use this parameter to display session information that matches the logout
or active task assignment wait time that you specify. Wait time refers to
the amount of time before attempting an explicit or implicit logout or
active task assignment after an unexpected connection termination or
connection reset.

``[-error-recovery-level <integer>]`` - Error Recovery Level::
Use this command to display session information that matches the error
recovery level that you specify.

``[-first-burst-length <integer>]`` - First Burst Length::
Use this parameter to display session information that matches the first
burst length that you specify. First burst length is the maximum amount of
unsolicited data in bytes that can be sent during the execution of a
single iSCSI packet. First burst length covers the total amount of
immediate data and the unsolicited data-out PDU. The first burst length
must not exceed the maximum burst length.

``[-immediate-data-enabled {true|false}]`` - Immediate Data::
Specifies if immediate data is supported. When immediate data is
supported, the initiator can send immediate data. If you enter this
command using the parameter without a value, it is set to true, and the
command displays all session information that supports immediate data. If
you provide a false value, the command displays all session information
that does not support immediate data.

``[-initiator-alias <text>]`` - Initiator Alias::
Use this parameter to display iSCSI session information that matches the
alias name of the initiator that you specify.

``[-initial-r2t-enabled {true|false}]`` - Initial R2T Enabled::
Specifies if the initiator supports Initial Ready to Transfer (R2T). R2T
is the mechanism that allows the target to request the initiator for
output data. If you enter this command using the parameter without a
value, it is set to true, and the command displays all session information
that supports initial R2T data. If you provide a false value, the command
displays all session information that does not support initial R2T data.

``[-initiator-name <text>]`` - Initiator Name::
Use this parameter to display the iSCSI session information that matches
the initiator name that you specify.
```

```
``[-isid <text>]`` - Initiator Session ID::
Use this parameter to display iSCSI session information that matches the
initiator session ID that you specify.

``[-max-burst-length <integer>]`` - Max Burst Length for Session::
Use this parameter to display iSCSI session information that matches the
maximum burst length that you specify. Maximum burst length is the maximum
iSCSI data payload in bytes for a data-in or solicited data-out sequence.

``[-max-connections <integer>]`` - Max Connections for Session::
Use this parameter to display iSCSI session information that matches the
maximum number of connections that you specify.

``[-max-outstanding-r2t <integer>]`` - Max Outstanding R2T for Session::
Use this parameter to display iSCSI session information that matches the
maximum number of outstanding R2T per task that you specify.

``[-session-type <iSCSI Session Type>]`` - Session Type::
Use this parameter to display iSCSI session information that matches the
session type that you specify.

``[-tpgroup-tag <integer>]`` - Target Portal Group Tag::
Use this parameter to display iSCSI session information that matches the
target portal group tag that you specify.

``[-connection-ids <integer>,...]`` - Active Connection IDs::
Use this parameter to display iSCSI session information that matches the
active connection IDs that you specify.

== Examples
```

cluster1::> vserver iscsi session show -vserver vs_1
Tpgroup Initiator Initiator
Vserver Name TSIH Name ISID Alias
--------- ------- ------ --------- --------- ----------
vs_1 tpgroup_1
2 iqn.1993-08.org.debian:01:fa752b8a5a3a
00:02:3d:01:00:00
initiator-alias
Displays session information for all sessions on Vserver vs_1.

```
== Related Links

* xref:{relative_path}vserver-iscsi-connection-show.html[vserver iscsi
connection show]
* xref:{relative_path}vserver-iscsi-session-parameter-show.html[vserver
iscsi session parameter show]
```

```
* xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify]




[[ID8438f2246bebfb4bd866569ce3a2846f]]
= vserver iscsi session shutdown
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Shut down a session on a node

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command forces a shutdown of all connections in a session. If you
want to shut down a single connection in a session, use the
xref:{relative_path}vserver-iscsi-connection-shutdown.html[vserver iscsi
connection shutdown] command.

== Parameters

``-vserver <Vserver Name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Specifies the Vserver.

``-tpgroup <text>`` - Target Portal Group [.advaced-priv]#(privilege:
advanced)#::
Specifies the target portal group that contains the session you want to
shutdown.

``-tsih <integer>`` - Target Session ID [.advaced-priv]#(privilege:
advanced)#::
Specifies the target session ID that you want to shut down.

== Examples
```

cluster1::*> vserver iscsi session shutdown -vserver vs_1 -tpgroup tpgroup_1 -tsih 2

```
Forces a session shutdown for target session ID 2 in tpgroup_1 in Vserver
vs_1 .
```

```
== Related Links

* xref:{relative_path}vserver-iscsi-connection-shutdown.html[vserver iscsi
connection shutdown]




[[ID88b463f665f2c971473d8b80283961ab]]
= vserver iscsi session parameter show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the parameters used to establish an iSCSI session

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command displays session parameter information. This command is
intended for troubleshooting performance problems.

Most of the parameters are read-only. However, some parameters can be
modified with the xref:{relative_path}vserver-iscsi-modify.html[vserver
iscsi modify] command.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display session information that matches the Vserver
name that you specify.

``[-tpgroup <text>]`` - Target Portal Group::
```

```
Use this parameter to display session information that matches the target
portal group name that you specify.

``[-tsih <integer>]`` - Target Session ID::
Use this parameter to display session information that matches the target
session ID that you specify.

``[-cmd-window-size <integer>]`` - Max Commands per Session::
Use this parameter to display session information that matches the command
window size that you specify.

``[-data-pdu-in-order {true|false}]`` - Data PDU in Order::
Use this parameter to display session information with the value of the
Protocol Data Units (PDU) in order flag you specify. This parameter
indicates if the data within a sequence can be in any order or must be in
sequence. If you enter this command without using this parameter, it is
set to true, and the command displays all session information that
supports PDUs in order. If you provide a false value, the command displays
all session information that does not support PDUs in order.

``[-data-sequence-in-order {true|false}]`` - Data Sequence in Order::
Use this parameter to display session information with the value of the
data sequence in order flag that you specify. If you enter this command
without using this parameter, it is set to true, and the command displays
all session information that supports data sequence. If you set the values
to false, the command displays all session information that does not
support data sequence.

``[-default-time-to-retain <integer>]`` - Default Time to Retain::
Use this parameter to display session information that matches the retain
time that you specify. This value specifies the amount of time before
active reassignment is possible after an unexpected connection termination
or a connection reset. A value of 0 means the connection task state is
immediately discarded by the target.

``[-default-time-to-wait <integer>]`` - Default Time to Wait::
Use this parameter to display session information that matches the logout
or active task assignment wait time that you specify. Wait time refers to
the amount of time before attempting an explicit or implicit logout or
active task assignment after an unexpected connection termination or
connection reset.

``[-error-recovery-level <integer>]`` - Error Recovery Level::
Use this command to display session information that matches the error
recovery level that you specify.
```

```
``[-first-burst-length <integer>]`` - First Burst Length::
Use this parameter to display session information that matches the first
burst length that you specify. First burst length is the maximum amount of
unsolicited data in bytes that can be sent during the execution of a
single iSCSI packet. First burst length covers the total amount of
immediate data and the unsolicited data-out PDU. The first burst length
must not exceed the maximum burst length.

``[-immediate-data-enabled {true|false}]`` - Immediate Data::
Use this parameter to display session information with the value of the
immediate data-enabled flag that you specify. If you enter this command
without using this parameter, it is set to true, and the command displays
all session information that supports immediate data. If you set the value
to false, the command displays all session information that does not
support immediate data.

``[-initial-r2t-enabled {true|false}]`` - Initial R2T Enabled::
Use this parameter to display session information with the value of the
R2T data-enabled flag that you specify. If you enter this command without
using this parameter, it is set to true, and the command displays all
session information that supports R2T data. If you set the value to false,
the command displays all session information that does not support R2T
data.

``[-initiator-alias <text>]`` - Initiator Alias::
Use this parameter to display iSCSI session information that matches the
initiator alias name you specify.

``[-initiator-name <text>]`` - Initiator Name::
Use this parameter to display iSCSI session information that matches the
initiator name you specify.

``[-isid <text>]`` - Initiator Session ID::
Use this parameter to display iSCSI session information that matches the
initiator session identifier you specify.

``[-max-burst-length <integer>]`` - Max Burst Length for Session::
Use this parameter to display iSCSI session information that matches the
maximum burst length that you specify. Maximum burst length is the maximum
iSCSI data payload in bytes for a data-in or solicited data-out sequence.

``[-max-connections <integer>]`` - Max Connections for Session::
Use this parameter to display iSCSI session information that matches the
maximum number of connections that you specify.

``[-max-outstanding-r2t <integer>]`` - Max Outstanding R2T for Session::
```

```
Use this parameter to display iSCSI session information that matches the
maximum number of outstanding R2T per task that you specify.

``[-session-type <iSCSI Session Type>]`` - Session Type::
Use this parameter to display iSCSI session information that matches the
session type you specify.

``[-tpgroup-tag <integer>]`` - Target Portal Group Tag::
Use this parameter to display iSCSI session information that matches the
target portal group tag you specify.

``[-initiator-mrdsl <integer>,...]`` - Initiator Max Recv Data Len::
Use this parameter to display iSCSI session information that matches the
initiator maximum receivable data segment length you specify. An iSCSI
initiator declares the maximum data segment length in bytes it can receive
in an iSCSI PDU during the iSCSI login phase.

``[-target-mrdsl <integer>,...]`` - Target Max Recv Data Len::
Use this parameter to display iSCSI session information that matches the
target maximum receivable data segment length you specify. An iSCSI target
declares the maximum data segment length in bytes it can receive in an
iSCSI PDU during the iSCSI login phase.

== Examples
```

cluster1::> iscsi session parameter show -vserver vs_1
Tpgroup Max Data PDU Data Seq Time 2 Time 2 Error Imm Initial
Vserver Name TSIH Conn In Order In Order Retain Wait Rec Lvl Data R2T
------- ------- ---- ---- -------- -------- ------ ------ ------- ----- -------
vs_1 vs_1.iscsi 6 1 true true 0 2 0 true false
vs_1 vs_1.iscsi 7 1 true true 0 2 0 true false
2 entries were displayed.

```
Lists iSCSI session parameters for Vserver vs_1.

== Related Links

* xref:{relative_path}vserver-iscsi-modify.html[vserver iscsi modify]




:leveloffset: -1


= vserver locks commands
```

```
:leveloffset: +1


[[IDf1cab5ada58f6e97010b8938ede03939]]
= vserver locks break
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Break file locks based on a set of criteria

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver locks break` command breaks one or more locks.

== Parameters

{ ``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver containing the lock.

``-volume <volume name>`` - Volume [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the volume containing the lock.

``[-lif <lif-name>]`` - Logical Interface [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the logical interface through which the lock was
established.

``[-lif-id <integer>]`` - Logical Interface ID [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the logical interface id through which the lock
was established.

``-path <text>`` - Object Path [.advaced-priv]#(privilege: advanced)#::
This parameter specifies a path to the lock.

| ``-lockid <UUID>`` - Lock UUID [.advaced-priv]#(privilege: advanced)#
}::
This parameter specifies the universally unique identifier (UUID) for the
lock. Queries and wildcard characters are not supported.
```

```
``-owner-id <text>`` - Owner ID [.advaced-priv]#(privilege: advanced)#::
This parameter specifies an owner ID for a lock. This parameter must be
used with the query notation { } exhibited in the second example.

``-protocol <lock protocol>`` - Lock Protocol [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the protocol that was used to establish a lock.
This parameter must be used with the query notation { } exhibited in the
second example.

``-client-address <IP Address>`` - Client Address [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies a client address associated with a lock. This
parameter must be used with the query notation { } exhibited in the second
example.

``-client-address-type {ipv4|ipv6|ipv6z}`` - Client Address Type
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the type of ip address a client used to create
its lock (ipv4, ipv6). This parameter must be used with the query notation
{ } exhibited in the second example.

``-flexcache-volume <text>`` - FlexCache Volume Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the FlexCache volume. This parameter
must be used with the query notation { } exhibited in the third example.

``-flexcache-vserver <text>`` - FlexCache Vserver Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the Vserver hosting the FlexCache
volume. This parameter must be used with the query notation { } exhibited
in the third example.

``-flexcache-cluster <text>`` - FlexCache Cluster Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the cluster hosting the FlexCache
volume. This parameter must be used with the query notation { } exhibited
in the third example.

== Examples

The following example breaks the locks on all objects on the Vserver named
vs0 in the volume named vol0, regardless of the paths to the locked
objects and the logical interface through which the locks were
established.
```

```
cluster1::*> vserver locks break -vserver vs0 -volume vol0 -path * -lif *
WARNING: Breaking file locks can cause applications to become unsynchronized
and may lead to data corruption.
Do you want to continue? {y|n}: y
1 entry was acted on.
```

```
The vserver locks break command can also be issued using a query on the
parameters available to the vserver locks show command. The following
example breaks all NLM protocol lock objects locked by the client at
address 12.34.56.78.
```

```
cluster1::*> vserver locks break { -protocol nlm -client-address 12.34.56.78 }
Warning: Breaking file locks can cause applications to become unsynchronized
and may lead to data corruption.
Do you want to continue? {y|n}: y
1 entry was acted on.
```

```
The following example breaks all FlexCache lock objects locked for
FlexCache volume "fc1" in Vserver "vs12".
```

```
cluster1::*> vserver locks break { -flexcache-volume fc1 -flexcache-vserver vs12 -flexcache-cluster cluster2 }
Warning: Breaking file locks can cause applications to become unsynchronized
and may lead to data corruption.
Do you want to continue? {y|n}: y
1 entry was acted on.
```

```
[[ID1ef2fc10d6e15a57573695f83ee4dccc]]
= vserver locks show-mirrored-open-count
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display mirrored open count

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver locks show-mirrored-open-count ` command displays number of
mirrored opens on node.

== 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.

| ``[-detailed  ]``::
This option displays detailed count of type of locks persisted on this
node.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node::
This parameter displays node name.

``[-mirrored-open-count <integer>]`` - Mirrored Open Count::
This parameter displays mirrored opens count on node.

``[-persistent-locks <integer>]`` - Persistent Opens::
This parameter displays count of opens persisted on node.

``[-mirrored-locks <integer>]`` - Mirrored Opens::
This parameter displays count of opens mirrored on node.

== Examples

The following example displays the mirrored opens count on node1.
```

cluster1::> vserver locks show-mirrored-open-count
Node Mirrored Open Count
---------- ----------
node1 40

```
The following example displays the mirrored opencount for HA pair cluster.
```

cluster1::> vserver locks show-mirrored-open-count
Node Mirrored Open Count
---------- ----------
node1 20
node2 30

```
[[IDe5cce4546d906ff262fc776f19b75d48]]
= vserver locks show
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display current list of locks

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver locks show` command displays information about locks. A lock
is a synchronization mechanism for enforcing limits on concurrent access
to files where many clients can be accessing the same file at the same
time. The command output depends on the parameter or parameters specified
with the command. If you do not specify any parameters, the command
displays the following information about locks:

* Vserver name
* Volume name
* Object path
* Logical interface name
* Lock protocol
* Lock type
* Client



NOTE: The `vserver locks show` command is also used to display FlexCache
specific locks. FlexCache locks are not stored on the FlexCache volume.
Instead all of the locks are stored on the origin of a FlexCache volume.
To view the FlexCache locks use the `vserver locks show` command on the
origin cluster.



== 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.

| ``[-smb-attrs  ]``::
If you specify the `-smb-attrs` parameter, the command displays
information related to SMB2 and higher.
```

```
| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

{ ``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information about
locks on the specified Vserver.

``[-volume <volume name>]`` - Volume::
If you specify this parameter, the command displays information about
locks on the specified volume.
+
NOTE: For the FlexCache locks, this parameter displays the origin of a
FlexCache volume name.
+


``[-lif <lif-name>]`` - Logical Interface::
If you specify this parameter, the command displays information about
locks established through the specified logical interface name.
+
NOTE: For the FlexCache locks, this parameter is unset if the FlexCache
volume is not local to the cluster.
+


``[-lif-id <integer>]`` - Logical Interface ID [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information about
locks established through the specified logical interface id.
+
NOTE: For the FlexCache locks, this parameter displays the logical
interface id of the FlexCache cluster.
+


``[-path <text>]`` - Object Path::
If you specify this parameter, the command displays information about
locks at the specified path name.

| ``[-lockid <UUID>]`` - Lock UUID }::
If you specify this parameter, the command displays information about the
lock with the specified universally unique identifier (UUID).

``[-is-constituent {true|false}]`` - Is Constituent Volume::
If you specify this parameter, the command displays information only about
volumes that either are or are not constituents of a FlexGroup, depending
```

```
on the value provided.

``[-protocol <lock protocol>]`` - Lock Protocol::
If you specify this parameter, the command displays information about
locks established through the specified protocol. Some of the valid
protocols are:
+
  * ``_cifs_`` : SMB locks
  * ``_nlm_`` : NFS3 locks
  * ``_nfsv4_`` : NFS4.0 locks
  * ``_nfsv4.1_`` : NFS4.1 locks
  * ``_crposix_`` : CrPosix locks for CREATE and LINK
  * ``_fcache_`` : Delegations for 7-mode destination FlexCache volumes



``[-type {byte-range|share-level|op-lock|delegation}]`` - Lock Type::
If you specify this parameter, the command displays information about
locks of the specified lock type. The four types of locks are:
+
  * Byte-range locks: Lock only a portion of a file.
  * Share locks: Represent opened files.
  * Opportunistic locks: Control client-side caching over SMB.
  * Delegations: Control client-side caching over NFSv4.



``[-node <nodename>]`` - Node Holding Lock State::
If you specify this parameter, the command displays information about all
locks on the specified node.

``[-lock-state <lock_state>]`` - Lock State::
If you specify this parameter, the command displays information about the
state of the lock. Some of the valid states are:
+
  * ``_granted_`` : The lock is established.
  * ``_revoking_`` : The server is currently coordinating with the client
to change the state of this lock.
  * ``_revoked_`` : The lock is undergoing revocation to be downgraded or
released.
  * ``_adjusted_`` : The lock is undergoing revocation to be replaced by a
lock equal to or weaker than the current lock.
  * ``_subsumed_`` : The lock is one of a set of locks that will replace a
lock that is being revoked.
  * ``_waiting_`` : The lock is waiting to be granted, because it
conflicts with another lock.
```

```
  * ``_denied_`` : The lock has been denied.
  * ``_timeout_`` : The lock was waiting and has now timed out.
  * ``_gone_`` : The lock is about to be released.
  * ``_unused_`` : The lock is allocated but has not been processed into
any state.



``[-bytelock-offset <integer>]`` - Bytelock Starting Offset::
If you specify this parameter, the command displays information about
bytelocks with the specified offset value. This is the index in the file
(in bytes) where the lock begins.

``[-bytelock-length <integer>]`` - Number of Bytes Locked::
If you specify this parameter, the command displays information about
bytelocks with the specified length. This is the number of bytes that are
locked by this particular lock.

``[-bytelock-mandatory {true|false}]`` - Bytelock is Mandatory::
If you specify this parameter, the command displays information only about
mandatory bytelocks. A mandatory bytelock enforces the requirement of byte
range locking on clients before accessing the associated range.

``[-bytelock-exclusive {true|false}]`` - Bytelock is Exclusive::
If you specify this parameter, the command displays information only about
exclusive bytelocks. When an exclusive bytelock is granted, no other
bytelock may be granted whose range overlaps it.

``[-bytelock-super {true|false}]`` - Bytelock is Superlock::
If you specify this parameter, the command displays information only about
super-bytelocks. When a super-bytelock is granted, all other locks on that
file are released, and no other operations will be allowed on that file.

``[-bytelock-soft {true|false}]`` - Bytelock is Soft::
If you specify this parameter, the command displays information only about
softened bytelocks. An NFSv4 bytelock might become softened if the
connection to the client is interrupted. Soft locks might be reclaimed if
the client reconnects. However if another client requests a lock that
conflicts with a soft lock, then the soft lock will be released.

``[-oplock-level {exclusive|level2|batch|null|read-batch}]`` - Oplock
Level::
If you specify this parameter, the command displays information about
locks with the specified oplock level. The oplock level determines which
operations the client may cache locally. Those operations include opening,
reading, writing, closing, and creating and destroying bytelocks on a
```

```
file. The five valid oplock levels are:
+
  * ``_batch_`` : The client may cache all operations on the file.
  * ``_exclusive_`` : The client may cache reads and writes on the file.
  * ``_read-batch_`` : The client may cache reads and opens on the file.
  * ``_level2_`` : The client may cache reads on the file.
  * ``_null_`` : The client may not cache any operations on the file.



``[-sharelock-mode <share lock mode>]`` - Shared Lock Access Mode::
If you specify this parameter, the command displays information about
locks with the specified sharelock mode. The parameter has two components
separated by a hyphen: the access mode followed by the share mode. The
access mode specifies which operations the client is allowed to perform on
the file. The share mode specifies which operations other clients are
disallowed to perform. The two modes are a combination of one or more of
these permissions:
+
  * ``_read_``
  * ``_write_``
  * ``_delete_``
  * ``_all_``
  * ``_none_``

+
For example, the sharelock mode ``_read_write-deny_delete_`` allows the
client to read and write the file, and disallows other clients to delete
the file. A special mode is ``_delete-on-close_`` , which specifies that
the server will delete the file as soon as it is closed.

``[-sharelock-soft {true|false}]`` - Shared Lock is Soft::
If you specify this parameter, the command displays information only about
softened sharelocks. A NFSv4 sharelock can become softened when the
connection to the client is interrupted. If the client reconnects, it
might reclaim the sharelock. However, if another client creates a
sharelock that conflicts with the softened sharelock, the softened
sharelock will be released.

``[-delegation-type {read|write}]`` - Delegation Type::
If you specify this parameter, the command displays information only about
locks with the specified delegation-type setting. The delegation type
determines which operations the client may cache locally. The two valid
delegation types are:
+
  * ``_read_`` : The client may cache reads on the file.
```

```
  * ``_write_`` : The client may cache reads and writes on the file.



``[-owner-id <text>]`` - Owner ID::
If you specify this parameter, the command displays information only about
locks with the specified owner ID. The owner ID is an opaque byte string
generated by the server for each file lock request.

``[-client-address <IP Address>]`` - Client Address::
If you specify this parameter, the command displays information only about
locks from the specified client IP address.

``[-client-address-type {ipv4|ipv6|ipv6z}]`` - Client Address Type::
If you specify this parameter, the command displays information only about
locks corresponding to a certain IP address type. Note that locks created
over the NFSv4 or NFSv4.1 protocol cannot have their address types
resolved. Valid options are:
+
  * ``_ipv4_`` : Clients operating over an IPv4 interface.
  * ``_ipv6_`` : Clients operating over an IPv6 interface.



``[-smb-open-type {none|durable|persistent}]`` - SMB Open Type::
If you specify this parameter, the command displays information only about
locks with the specified SMB open type. Valid open types are
+
  * ``_durable_`` : Durability is a feature of SMB2. A durable lock might
become "disconnected" if the connection between the client and server is
disrupted. A disconnected durable lock might be reconnected if the
connection is reestablished.
  * ``_persistent_`` : Persistence is a feature of SMB3. Persistent locks
can become disconnected and later reconnected, like durable locks.
Persistent locks are used to facilitate continuously available shares.
  * ``_none_`` : The lock is neither durable nor persistent.



``[-smb-connect-state <Lock Connect State>]`` - SMB Connect State::
If you specify this parameter, the command displays information only about
locks with the specified SMB connection state. Some of the valid states
are:
+
  * ``_connected_`` : This is the normal state of a SMB lock when the
server and client are connected.
  * ``_disconnected_`` : If a lock is durable or persistent, it might
```

```
become disconnected if the connection between the server and its client is
interrupted. Disconnected locks may later be reconnected if the connection
is reestablished.



``[-smb-expiration-time <integer>]`` - SMB Expiration Time (Secs)::
If you specify this parameter, the command displays information only about
locks with the specified SMB lock expiration time. When a lock is
disconnected, `-smb-expiration-time` shows the time remaining until the
lock expires. The server releases the lock after it expires.

``[-smb-open-group-id <text>]`` - SMB Open Group ID::
If you specify this parameter, the command displays information only about
locks with the specified SMB open group identifier. This is an opaque byte
string provided by the client as the SMB lease key when the lock is first
established.

``[-is-flexcache-lock {true|false}]`` - Is FlexCache Lock::
If you specify this parameter, the command displays information only about
locks with the specified value.

``[-flexcache-volume <text>]`` - FlexCache Volume Name::
If you specify this parameter, the command displays information only about
locks with the specified FlexCache volume.

``[-flexcache-vserver <text>]`` - FlexCache Vserver Name::
If you specify this parameter, the command displays information only about
locks with the specified Vserver hosting a FlexCache volume.

``[-flexcache-vserver-uuid <UUID>]`` - FlexCache Vserver UUID::
If you specify this parameter, the command displays information only about
locks with the specified UUID of a Vserver hosting a FlexCache volume.

``[-flexcache-volume-msid <integer>]`` - FlexCache Volume MSID::
If you specify this parameter, the command displays information only about
locks with the specified FlexCache volume MSID

``[-flexcache-cluster <text>]`` - FlexCache Cluster Name::
If you specify this parameter, the command displays information only about
locks with the specified cluster hosting a FlexCache volume.

``[-flexcache-lock-authority <text>]`` - FlexCache lock authority::
If you specify this parameter, the command displays information only about
locks with the specified authority granted to a FlexCache volume.

== Examples
```

```
The following example displays default information about all locks:
```

cluster1::> vserver locks show

Vserver: vs0
Volume Object Path LIF Protocol Lock Type Client
-------- ------------------------ ----------- --------- ----------- ----------
vol1 /vol1/notes.txt node1_data1
cifs share-level 192.168.1.5
Sharelock Mode: read_write-deny_delete
op-lock 192.168.1.5
Oplock Level: read-batch
/vol1/notes1.txt node1_data1
cifs share-level 192.168.1.5
Sharelock Mode: read_write-deny_delete
op-lock 192.168.1.5
Oplock Level: batch
/vol1 node1_data2
cifs share-level 192.168.1.5
Sharelock Mode: read-deny_delete
/vol1/notes.txt node1_data2
cifs share-level 192.168.1.5
Sharelock Mode: read_write-deny_delete
op-lock 192.168.1.5
Oplock Level: read-batch
7 entries were displayed.

```
The following example displays the SMB related information about all
locks:
```

cluster1::> vserver locks show -smb-attrs

Vserver: vs0
Volume Object Path LIF Protocol Lock Type Client
-------- ------------------------ ----------- --------- ----------- ----------
vol1 /vol1/notes.txt node1_data1
cifs share-level 192.168.1.5
Sharelock Mode: read_write-deny_delete
Open Type: durable Connect State: connected Expiration Time (Secs): -
Open Group ID: 625e2ff46ee5df1194ba0050569d37047058909c00000000873d210700000000
op-lock 192.168.1.5
Oplock Level: read-batch
Open Type: - Connect State: connected Expiration Time (Secs): -
Open Group ID: 625e2ff46ee5df1194ba0050569d37047058909c00000000873d210700000000
/vol1/notes1.txt node1_data1
cifs share-level 192.168.1.5
Sharelock Mode: read_write-deny_delete
Open Type: durable Connect State: connected Expiration Time (Secs): -
Open Group ID: 625e2ff46ee5df1194ba0050569d370440fc8891000000005a3f210700000000
op-lock 192.168.1.5

Oplock Level: batch
Open Type: - Connect State: connected Expiration Time (Secs): -
Open Group ID: 625e2ff46ee5df1194ba0050569d370440fc8891000000005a3f210700000000
/vol1 node1_data2
cifs share-level 192.168.1.5
Sharelock Mode: read-deny_delete
Open Type: none Connect State: connected Expiration Time (Secs): -
Open Group ID: -
/vol1/notes.txt node1_data2
cifs share-level 192.168.1.5
Sharelock Mode: read_write-deny_delete
Open Type: durable Connect State: connected Expiration Time (Secs): -
Open Group ID: 625e2ff46ee5df1194ba0050569d370408e08d9c00000000da40210700000000
op-lock 192.168.1.5
Oplock Level: read-batch
Open Type: - Connect State: connected Expiration Time (Secs): -
Open Group ID: 625e2ff46ee5df1194ba0050569d370408e08d9c00000000da40210700000000

7 entries were displayed.

```
The following example displays default information about all locks in list
form:
```

cluster1::> vserver locks show -instance
Vserver: vs0
Volume: vol1
Logical Interface: node1_data1
Object Path: /vol1/notes.txt
Lock UUID: 447db184-f801-11df-8bb5-00a098000e34
Lock Protocol: cifs
Lock Type: share-level
Node Holding Lock State: node1
Lock State: granted
Bytelock Starting Offset: -
Number of Bytes Locked: -
Bytelock is Mandatory: -
Bytelock is Exclusive: -
Bytelock is Superlock: -
Bytelock is Soft: -
Oplock Level: -
Shared Lock Access Mode: read_write-deny_delete
Shared Lock is Soft: false
Delegation Type: -
Client Address: 192.168.1.5
Client Address Type: ipv4
SMB Open Type: durable
SMB Connect State: connected
SMB Expiration Time (Secs): -
SMB Open Group ID: 625e2ff46ee5df1194ba0050569d37047058909c00000000873d2107000000004
Vserver: vs0
Volume: vol1
Logical Interface: node1_data1

Object Path: /vol1/notes.txt
Lock UUID: 447db185-f801-11df-8bb5-00a098000e34
Lock Protocol: cifs
Lock Type: op-lock
Node Holding Lock State: node1
Lock State: granted
Bytelock Starting Offset: -
Number of Bytes Locked: -
Bytelock is Mandatory: -
Bytelock is Exclusive: -
Bytelock is Superlock: -
Bytelock is Soft: -
Oplock Level: read-batch
Shared Lock Access Mode: -
Shared Lock is Soft: -
Delegation Type: -
Client Address: 192.168.1.5
Client Address Type: ipv4
SMB Open Type: -
SMB Connect State: connected
SMB Expiration Time (Secs): -
SMB Open Group ID: 625e2ff46ee5df1194ba0050569d37047058909c00000000873d210700000000
Vserver: vs0
Volume: vol1
Logical Interface: node1_data1
Object Path: /vol1/notes1.txt
Lock UUID: 48cee334-f801-11df-8bb5-00a098000e34
Lock Protocol: cifs
Lock Type: share-level
Node Holding Lock State: node1
Lock State: granted
Bytelock Starting Offset: -
Number of Bytes Locked: -
Bytelock is Mandatory: -
Bytelock is Exclusive: -
Bytelock is Superlock: -
Bytelock is Soft: -
Oplock Level: -
Shared Lock Access Mode: read_write-deny_delete
Shared Lock is Soft: false
Delegation Type: -
Client Address: 192.168.1.5
Client Address Type: ipv4
SMB Open Type: durable
SMB Connect State: connected
SMB Expiration Time (Secs): -
SMB Open Group ID: 625e2ff46ee5df1194ba0050569d370440fc8891000000005a3f210700000000
3 entries were displayed.

```
The following example displays the FlexCache locks:
```

cluster1::> vserver locks show

```
Vserver: vs34
Volume Object Path LIF Protocol Lock Type Client
-------- ------------------------ ----------- --------- ----------- ----------
origin /origin/file1 data1 nlm byte-range 10.235.224.139
Bytelock Offset(Length): 0 (18446744073709551615)
/origin/file3 data1 nlm byte-range 10.234.133.121
Bytelock Offset(Length): 0 (18446744073709551615)
FlexCache Lock: true
/origin/file2 - nlm byte-range 10.234.189.144
Bytelock Offset(Length): 0 (18446744073709551615)
FlexCache Lock: true
3 entries were displayed.
cluster1::> vserver locks show -fields flexcache-vserver, flexcache-volume, flexcache-cluster

vserver volume lif path flexcache-volume flexcache-vserver flexcache-cluster
------- ------ ----- ------------ ---------------- ----------------- -----------------
vs34 origin data1 /origin/file1 - - -
vs34 origin data1 /origin/file3 local_fc vs34 cluster2
vs34 origin - /origin/file2 fc1 vs12 cluster1
3 entries were displayed.
```

```
The following example displays default information about FlexCache and
origin of FlexCache locks in list form:
```

```
cluster1::> vserver locks show -instance
Vserver: vs34
Volume: origin
Logical Interface: data1
Object Path: /origin/file1
Lock UUID: 173d10d6-f64c-4983-a2c2-9db0bc4f5c05
Is Constituent Volume: false
Lock Protocol: nlm
Lock Type: byte-range
Node Holding Lock State: node1
Lock State: granted
Bytelock Starting Offset: 0
Number of Bytes Locked: 18446744073709551615
Bytelock is Mandatory: false
Bytelock is Exclusive: false
Bytelock is Superlock: false
Bytelock is Soft: false
Oplock Level: -
Shared Lock Access Mode: -
Shared Lock is Soft: -
Delegation Type: -
Owner ID: 31-113439407363737072303435323636363030312e67646c2e656e676c61622e6e65746170
Client Address: 10.235.224.139
Client Address Type: ipv4
SMB Open Type: -
SMB Connect State: -
SMB Expiration Time (Secs): -
SMB Open Group ID: -
```

Is FlexCache Lock: -
FlexCache Volume Name: -
FlexCache Vserver Name: -
FlexCache Vserver UUID: -
FlexCache Volume MSID: -
FlexCache Cluster Name: -
Vserver: vs34
Volume: origin
Logical Interface: data1
Object Path: /origin/file3
Lock UUID: e3bf0c78-7371-41b6-9eff-4fa81d64ca08
Is Constituent Volume: false
Lock Protocol: nlm
Lock Type: byte-range
Node Holding Lock State: node1
Lock State: granted
Bytelock Starting Offset: 0
Number of Bytes Locked: 18446744073709551615
Bytelock is Mandatory: false
Bytelock is Exclusive: false
Bytelock is Superlock: false
Bytelock is Soft: false
Oplock Level: -
Shared Lock Access Mode: -
Shared Lock is Soft: -
Delegation Type: -
Owner ID: 25-1133374073637370723034393434333130303 1
Client Address: 10.234.133.121
Client Address Type: ipv4
SMB Open Type: -
SMB Connect State: -
SMB Expiration Time (Secs): -
SMB Open Group ID: -
Is FlexCache Lock: true
FlexCache Volume Name: local_fc
FlexCache Vserver Name: vs34
FlexCache Vserver UUID: 388e3004-7520-11e8-8789-005056b47572
FlexCache Volume MSID: 2160150696
FlexCache Cluster Name: cluster2
Vserver: vs34
Volume: origin
Logical Interface: -
Object Path: /origin/file2
Lock UUID: c420b002-d78c-4157-94d5-55c186ef4df3
Is Constituent Volume: false
Lock Protocol: nlm
Lock Type: byte-range
Node Holding Lock State: node1
Lock State: granted
Bytelock Starting Offset: 0
Number of Bytes Locked: 18446744073709551615
Bytelock is Mandatory: false
Bytelock is Exclusive: false
Bytelock is Superlock: false

Bytelock is Soft: false
Oplock Level: -
Shared Lock Access Mode: -
Shared Lock is Soft: -
Delegation Type: -
Owner ID: 32-1135304073637370723034353236363430303 1
Client Address: 10.234.189.144
Client Address Type: ipv4
SMB Open Type: -
SMB Connect State: -
SMB Expiration Time (Secs): -
SMB Open Group ID: -
Is FlexCache Lock: true
FlexCache Volume Name: fc1
FlexCache Vserver Name: vs12
FlexCache Vserver UUID: 5a943dd9-7520-11e8-b5e7-005056b47786
FlexCache Volume MSID: 2150871844
FlexCache Cluster Name: cluster1
3 entries were displayed.

```
:leveloffset: -1


= vserver migrate commands

:leveloffset: +1


[[ID1451315df551ad6add7459f51851cd43]]
= vserver migrate abort
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Abort a Vserver migrate operation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command aborts a Vserver migrate operation. It is normally run on the
destination cluster of the migrate operation. If the destination cluster
cannot be accessed, the abort command can be run on the source cluster to
cleanup the migrate operation on the source cluster. The migrate operation
must be in migrate-paused, or in a failed migrate status, to be aborted.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver Name::
Name of the Vserver to be aborted.

== Examples

cluster1::> vserver migrate abort -vserver test




[[ID1bb5cc734bb976fdf0e2e008db978037]]
= vserver migrate cutover
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Perform Cutover of the migrate operation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command performs a cutover of the Vserver from the source cluster to
the destination cluster. It must be run on the destination cluster of the
Vserver migrate.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Name of the Vserver which is being migrated.

== Examples
```

cluster1::> vserver migrate cutover -vserver test

```
[[ID1f30c6336eac76c082f9e0611e2ebc35]]
= vserver migrate pause
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Pause a Vserver migrate operation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command pauses a Vserver migrate operation. Both data transfer and
configuration replication are stopped. It must be run on the destination
cluster of the Vserver migrate operation.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Name of the Vserver whose migration operation will be paused.

== Examples

cluster1::> vserver migrate pause -vserver test




[[ID029e1a0618f493b8542cdcb8478ee184]]
= vserver migrate resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume a migrate operation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command resumes a Vserver migrate operation. It must be run on the
destination cluster of the Vserver migrate operation. The migrate
operation, which either previously failed or was paused in order to
prioritize other cluster operations, can be resumed.

== Parameters
```

```
``-vserver <vserver name>`` - Vserver Name::
Name of the Vserver being migrated.

``[-auto-cutover {true|false}]`` - Automatically cutover when ready::
This parameter specifies if the Vserver migrate operation should cutover
automatically when ready. If this parameter is not specified, migrate
resume will use the value from the previous migrate start or resume
operation.

``[-auto-source-cleanup {true|false}]`` - Automatically cleanup the source
vserver after cutover::
This parameter specifies if the Vserver migrate operation should cleanup
the source Vserver automatically. Setting the parameter requires advanced
privilege. If this parameter is not specified, migrate resume will use the
value from the previous migrate start or resume operation.

``[-throttle <throttleType>]`` - Throttle for Migrate Transfers(Kbs)::
 This parameter specifies the throttle value to limit the network
bandwidth used for the migrate transfers. It sets the maximum rate (in
Kbytes/sec) at which data can be transferred during the operation. The
throttle value is applied individually to each volume of the migrating
Vserver. To fully use the network bandwidth available, set the throttle
value to unlimited or 0. If this parameter is not specified, migrate will
use the throttle value from the previous migrate start or resume
operation. The minimum throttle value is four Kbytes/sec, so if you
specify a throttle value between 1 and 4, it will be treated as if you
specified 4.

== Examples

cluster1::> vserver migrate resume -vserver test -auto-cutover false
cluster1::> vserver migrate resume -vserver test -auto-cutover false
-throttle 12500




[[ID8d2062344a3f9167430607976f7e5c7a]]
= vserver migrate show-volume-fix-details
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
List volumes and commands to run post migrate
```

```
*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command displays the list of volumes whose properties failed to get
transferred to the destination Vserver of the Vserver migrate post
cutover.

For each of these volumes, this command also displays the
xref:{relative_path}volume-modify.html[volume modify] command which you
must run to transfer the failed volume properties manually.

== Parameters

``[-fields <fieldname>,...]``::
This specifies the fields that must be displayed.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is specified, the command displays the volume details
and properties that failed to get transferred during the migrate.

``[-vserver-uuid <UUID>]`` - Vserver-UUID [.advaced-priv]#(privilege:
advanced)#::
The UUID of the Vserver which is migrating.

``[-volume <volume name>]`` - Volume Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays the volume details
and properties that match the specified volume.

``[-command <text>]`` - Volume Modify Command [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays the volume details
and properties that match the specified `Volume Modify` command.

== Examples
```

clusterB::> vserver migrate show-volume-fix-details -vserver-uuid 7a08a236-f6b3-11eb-98a5-005056a78af4
Vserver Volume Command
-------- ------------ ---
7a08a236-f6b3-11eb-98a5-005056a78af4
vs4_root
volume modify -vserver vsnew -volume vs4_root -min-autosize 20971520B -autosize-grow-threshold-percent 85% -autosize-shrink-threshold-percent 50% -autosize-mode off -files 566 -space-slo none -space-guarantee volume -fractional-reserve 100% -min-readahead false -atime-update true -snapdir-access true -percent -snapshot-space 5% -space-mgmt-try-first volume_grow -read-realloc space-optimized -is-space-enforcement

-logical false -activity-tracking-state off -snap-autodelete-enabled false -snap-autodelete-commitment try -snap -autodelete-defer-delete user_created -snap-autodelete-delete-order oldest_first -snap-autodelete-prefix "(not specified)" -snap-autodelete-target-free-space 20% -snap-autodelete-trigger volume -snap-autodelete-destroy -list none

```
== Related Links

* xref:{relative_path}volume-modify.html[volume modify]




[[ID929c4e2be97c0ce2e0caf3fb813bb965]]
= vserver migrate show-volume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display status of volumes in a migrating Vservers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays status of all the volumes in a migrating Vservers.
It must be run on the destination cluster of the Vserver migrate.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed volume
status information.

``[-vserver <vserver name>]`` - Vserver Name::
Name of the Vserver which is migrating.

``[-volume <volume name>]`` - Volume Name::
If this parameter is specified, the command displays detailed volume
status information that matches the specified volume.

``[-volume-dsid <integer>]`` - Volume DSID::
```

```
If this parameter is specified, the command displays the details of volume
status that matches the specified DSID.

``[-volume-msid <integer>]`` - Volume MSID::
If this parameter is specified, the command displays the details of volume
status that matches the specified MSID.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
If this parameter is specified, the command displays the details of volume
status that matches the specified Vserver UUID.

``[-node <nodename>]`` - Node::
If this parameter is specified, the command displays the details of volume
status that matches the specified node.

``[-volume-state {online|restricted|offline|force-online|force-
offline|mixed}]`` - State of the Volume::
If this parameter is specified, the command displays the details of volume
status that matches the specified volume state.

``[-transfer-status <mirror status>]`` - Status of Transfer::
If this parameter is specified, the command displays the details of volume
status that matches the specified transfer status.

``[-healthy {true|false}]`` - Healthy::
If this parameter is specified, the command displays the details of volume
status that matches the specified health status

``[-errors <text>,...]`` - Errors in Volume Operation::
If this parameter is specified, the command displays the details of volume
status that matches the specified errors.

== Examples

cluster1::> vserver migrate show-volume -vserver test




[[ID1f098bd1cb031ac5e38a95dcafc863f4]]
= vserver migrate show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Display status of migrating Vservers

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays information about the migrating vserver.

== 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.

| ``[-migrate-status-details  ]``::
If this parameter is specified, the command displays the following
information about Vserver migrate operation.
+
  * Vserver
  * Migrate status
  * Status details



| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver Name::
Name of the Vserver which is migrating.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver UUID.

``[-destination-cluster <Cluster name>]`` - Destination Cluster Name::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified destination cluster.

``[-source-cluster <Cluster name>]`` - Source Cluster Name::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified source cluster.

``[-migrate-status {precheck-started|transferring|ready-for-
cutover|cutover-triggered|cutover-started|cutover-complete|migrate-
```

```
paused|migrate-complete|migrate-complete-with-warnings|migrate-
failed|cleanup|cleanup-failed|manual-cleanup|destination-cleaned-
up|migrate-pausing|cleanup-pausing|post-cutover|post-cutover-failed|ready-
for-source-cleanup|setup-configuration|setup-configuration-failed|migrate-
pause-failed|migrate-aborting|migrate-abort-failed|migrate-aborted}]`` -
Vserver migrate status::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate status.

``[-start-time <MM/DD/YYYY HH:MM:SS>]`` - Migrate start time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate start time.

``[-end-time <MM/DD/YYYY HH:MM:SS>]`` - Migrate operation finish time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate end time.

``[-last-pause-time <MM/DD/YYYY HH:MM:SS>]`` - Last migrate pause time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate last paused time.

``[-last-resume-time <MM/DD/YYYY HH:MM:SS>]`` - Last migrate resume time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate last resume time.

``[-last-restart-time <MM/DD/YYYY HH:MM:SS>]`` - Last migrate restart
time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate last rollback time.

``[-cutover-precommit-start-time <MM/DD/YYYY HH:MM:SS>]`` - Cutover
precommit start time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate cutover trigger
time.

``[-cutover-commit-start-time <MM/DD/YYYY HH:MM:SS>]`` - Cutover commit
start time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate cutover start time.

``[-cutover-complete-time <MM/DD/YYYY HH:MM:SS>]`` - Cutover complete
time::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate cutover completion
time.
```

```
``[-restart-count <integer>]`` - Restart Count::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified rollback count.

``[-status-details <text>,...]`` - Errors and Warnings During Migrate::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate status details.

``[-ipspace <IPspace>]`` - Destination cluster IPspace Name for vserver::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified ipspace.

``[-aggr-list <aggregate name>,...]`` - Aggregate list for creating the
volumes::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified aggregate list that are assigned
for Vserver to use.

``[-migrate-vserver-type {migrate-source|migrate-destination}]`` -
Identify if the Vserver is migrate source or destination::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver migrate type.

``[-owning-node <nodename>]`` - Name of node on which migrate is
orchestrated::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified node that owns Vserver migrate.

``[-stream-uuid <UUID>]`` - Stream Uuid::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified stream UUID.

``[-is-past-point-of-no-return {true|false}]`` - Indicate point of no
return for migrate::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified past point of no return flag.

``[-last-failed-migrate-status {precheck-started|transferring|ready-for-
cutover|cutover-triggered|cutover-started|cutover-complete|migrate-
paused|migrate-complete|migrate-complete-with-warnings|migrate-
failed|cleanup|cleanup-failed|manual-cleanup|destination-cleaned-
up|migrate-pausing|cleanup-pausing|post-cutover|post-cutover-failed|ready-
for-source-cleanup|setup-configuration|setup-configuration-failed|migrate-
pause-failed|migrate-aborting|migrate-abort-failed|migrate-aborted}]`` -
Migrate status when the migrate aborted or failed::
If this parameter is specified, the command displays Vserver migrate
```

```
information that matches the specified last failed Vserver migrate status.

``[-current-migrate-operation
{none|start|resume|pause|cleanup|cutover|abort}]`` - Current Migrate
operation::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified current Vserver migrate operation.

``[-last-migrate-operation
{none|start|resume|pause|cleanup|cutover|abort}]`` - Last Migrate
operation::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified last Vserver migrate operation.

``[-local-vserver-id <integer>]`` - Vserver ID::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified local Vserver ID.

``[-group-id <integer>]`` - Group ID::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified group ID.

``[-partner-vserver-id <integer>]`` - Partner Vserver ID::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified Vserver ID of a partner node.

``[-partner-group-id <integer>]`` - Partner group ID::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified group ID of a partner node.

``[-auto-cutover {true|false}]`` - Automatic cutover::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified auto cutover flag.

``[-skip-performance-check {true|false}]`` - Skip checking iops
requirement of volume on destination aggregates::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified skip performance flag.

``[-migration-uuid <UUID>]`` - Migration UUID::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified migration UUID.

``[-migration-failures <SmdbError>,...]`` - List of migration failures::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified migration failures.
```

```
``[-client-outage-window <integer>]`` - Window duration during cutover
when clients are not served::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified client outage window.

``[-auto-source-cleanup {true|false}]`` - Automatic source cleanup
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified auto source cleanup flag.

``[-throttle <throttleType>]`` - Throttle for Migrate Transfers(Kbs)::
If this parameter is specified, the command displays Vserver migrate
information that matches the specified throttle value.

== Examples

cluster1::> vserver migrate show -vserver test




[[ID4fd25fd10ce19b54053ea71d484d53ac]]
= vserver migrate source-cleanup
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Perform vserver cleanup on the source cluster

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

This command cleans up the migrate source Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Name of the Vserver to be deleted on the source cluster.

``[-force {true|false}]`` - Force option to override veto checks
[.advaced-priv]#(privilege: advanced)#::
This parameter is used to forcibly delete the Vserver.
```

```
== Examples

cluster1::> vserver migrate source-cleanup -vserver test




[[ID917577317a5fd08fedad7dfaa79a21e6]]
= vserver migrate start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start the Vserver migrate operation

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command starts the migration of a Vserver from one cluster to
another. This has to be run on the destination cluster, i.e. the cluster
where the Vserver is intended to migrate. The source cluster from where
the Vserver is to be migrated is specified in the command.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Name of the Vserver which needs to be migrated.

``-source-cluster <Cluster name>`` - Source Cluster Name::
Name of the source cluster.

``[-check-only {true|false}]`` - Check if migrate can be done::
Runs the prechecks and tells if the Vserver migrate operation can be
started or not.

``[-ipspace <IPspace>]`` - Destination cluster IPspace Name for vserver::
Name of the IPspace in the destination cluster.

``[-aggr-list <aggregate name>,...]`` - Aggregate list::
Provide the list of aggregates where the volumes will be created in the
destination cluster.
```

```
``[-auto-cutover {true|false}]`` - Automatically cutover when ready::
This parameter specifies if the Vserver migrate operation should cutover
automatically when ready. The default setting is true.

``[-auto-source-cleanup {true|false}]`` - Automatically cleanup the source
vserver after cutover::
This parameter is to specify if the Vserver migrate operation should
cleanup the source Vserver automatically. Setting the parameter requires
diagnostic privilege. The default setting is true.

``[-throttle <throttleType>]`` - Throttle for Migrate Transfers(Kbs)::
This parameter specifies the throttle value to limit the network bandwidth
used for the migrate transfers. It sets the maximum rate (in Kbytes/sec)
at which data can be transferred during the operation. The throttle value
is applied individually to each volume of the migrating Vserver. To fully
use the network bandwidth available, set the throttle value to unlimited
or 0. The default setting is unlimited. The minimum throttle value is four
Kbytes/sec, so if you specify a throttle value between 1 and 4, it will be
treated as if you specified 4.

== Examples

cluster1::> vserver migrate start -vserver test -source-cluster cluster-22
-ipspace ips1 -check-only true -aggr-list aggr1,aggr2 -auto-cutover false
-throttle 12500




:leveloffset: -1


= vserver name-mapping commands

:leveloffset: +1


[[IDa186093c88dadb8a377fdf2b10f338f8]]
= vserver name-mapping create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a name mapping
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver name-mapping create` command creates a name mapping. Name
mappings are applied in the order in which they occur in the priority
list; for example, a name mapping that occurs at position 2 in the
priority list is applied before a name mapping that occurs at position 3.
Each mapping direction (Kerberos-to-UNIX, Windows-to-UNIX, UNIX-to-
Windows, S3-to-UNIX and S3-to-Windows) has its own priority list. ONTAP
prevents you from creating two name mappings with the same pattern.

Patterns can be expressed as POSIX regular expressions. For information
about regular expressions, see the UNIX reference page for ``_regex(7)_``
.

Each Vserver can have up to 12500 name mappings in each direction.

NOTE: If you are using the CLI, you must delimit all regular expressions
with double quotation marks (\"). For instance, to enter the regular
expression ``_(.+)_`` in the CLI, type ``_"(.+)"_`` at the command prompt.
To add a "?" to the expression, press ESC followed by the "?".



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to create the name
mapping.

``-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}`` - Direction::
This parameter specifies the direction of the name mapping. Possible
values are ``_krb-unix_`` for a Kerberos-to-UNIX name mapping, ``_win-
unix_`` for a Windows-to-UNIX name mapping, ``_unix-win_`` for a UNIX-to-
Windows name mapping, ``_s3-unix_`` for a S3-to-UNIX name mapping and
``_s3-win_`` for a S3-to-Windows name mapping.

``-position <integer>`` - Position::
This parameter specifies the name mapping's position in the priority list.
Specify the position as a positive integer.
+
NOTE: If you want to create a new name mapping at a position that is
already occupied in the priority list, use the
xref:{relative_path}vserver-name-mapping-insert.html[vserver name-mapping
insert] command instead of the `vserver name-mapping create` command.
```

```
+


``-pattern <text>`` - Pattern::
This parameter specifies the pattern you want to match. Refer to the
command description section for details. The pattern can be up to 256
characters in length.

``-replacement <text>`` - Replacement::
This parameter specifies the replacement pattern. The replacement pattern
can be up to 256 characters in length.

{ ``[-address <IP Address/Mask>]`` - IP Address with Subnet Mask::
This optional parameter specifies the IP address that can be used to match
the client's workstation IP address with the pattern.

| ``[-hostname <text>]`` - Hostname }::
This optional parameter specifies the hostname that can be used to match
the corresponding client's workstation IP address with the list of IP
addresses with the pattern.

== Examples

The following example creates a name mapping on a Vserver named vs1. The
mapping is from UNIX to Windows at position 5 in the priority list. The
mapping maps the pattern cifs to the replacement EXAMPLE\Domain Users.
```

cluster1::> vserver name-mapping create -vserver vs1 -direction unix-win -position 5 -pattern jane_doe -replacement contoso\\jdoe -address 10.238.33.245/24
cluster1::> vserver name-mapping create -vserver vs1 -direction unix-win -position 6 -pattern john_smith -replacement contoso\\jsmith -hostname google.com

```
== Related Links

* xref:{relative_path}vserver-name-mapping-insert.html[vserver name-
mapping insert]




[[ID5906fb4f2952faf6ec115ac9ddb71289]]
= vserver name-mapping delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a name mapping

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver name-mapping delete` command deletes a name mapping.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver from which you want to delete the
name mapping.

``-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}`` - Direction::
This parameter specifies the direction of the name mapping that you want
to delete.

``-position <integer>`` - Position::
This parameter specifies the position of the name mapping that you want to
delete. Specify the position as a positive integer.

== Examples

The following example deletes a name mapping on a Vserver named vs1. The
name mapping is from UNIX to Windows and is at position 5.
```

cluster1::> vserver name-mapping delete -vserver vs1 -direction unix-win -position 5

```
[[IDa8fc79979a3d9a67f1d50713f8bb3e05]]
= vserver name-mapping insert
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a name mapping at a specified position

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver name-mapping insert` command creates a name mapping at a
specified position in the priority list. The command rearranges the list
as needed to accommodate the new entry. For instance, if you have a
priority list of five mappings and insert a new mapping at position 3, the
mapping previously at position 3 is moved to position 4, the mapping
previously at position 4 is moved to position 5, and the mapping
previously at position 5 is moved to position 6. Each mapping direction
(Kerberos-to-UNIX, Windows-to-UNIX, UNIX-to-Windows, S3-to-UNIX and S3-to-
Windows) has its own priority list.

You can specify patterns as POSIX regular expressions. For information
about regular expressions, see the UNIX reference page for ``_regex(7)_``
.

Each Vserver can have up to 12500 name mappings in each direction.

NOTE: If you are using the CLI, you must delimit all regular expressions
with double quotation marks ("). For instance, to enter the regular
expression ``_(.+)_`` in the CLI, type ``_"(.+)"_`` at the command prompt.
To add a "?" to the expression, press ESC followed by the "?".



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to create the name
mapping.

``-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}`` - Direction::
This parameter specifies the direction of the name mapping. Possible
values are ``_krb-unix_`` for a Kerberos-to-UNIX name mapping, ``_win-
```

```
unix_`` for a Windows-to-UNIX name mapping, ``_unix-win_`` for a UNIX-to-
Windows name mapping, ``_s3-unix_`` for a S3-to-UNIX name mapping and
``_s3-win_`` for a S3-to-Windows name mapping.

``-position <integer>`` - Position::
This parameter specifies the position in the priority list at which you
want to insert the new name mapping. Specify a position as a positive
integer.

``-pattern <text>`` - Pattern::
This parameter specifies the pattern you want to match. Refer to the
command description section for details. The pattern can be up to 256
characters in length.

``-replacement <text>`` - Replacement::
This parameter specifies the replacement pattern. The replacement pattern
can be up to 256 characters in length.

{ ``[-address <IP Address/Mask>]`` - IP Address with Subnet Mask::
This optional parameter specifies the IP address that can be used to match
the client's workstation IP address with the pattern.

| ``[-hostname <text>]`` - Hostname }::
This optional parameter specifies the hostname that can be used to match
the corresponding client's workstation IP address with the list of IP
addresses with the pattern.

== Examples

The following example creates a name mapping on a Vserver named vs1. It is
a user mapping from Kerberos to UNIX. It is inserted into the priority
list at position 2. The name mapping maps any principal in the Kerberos
realm SEC.EXAMPLE.COM to the UNIX user name corresponding to the
principal's base name with any instance names removed; for example,
tom/admin@SEC.EXAMPLE.COM is mapped to tom.
```

cluster1::> vserver name-mapping insert -vserver vs1 -direction krb-unix -position 2 -pattern "(
)(/[^@])?@SEC.EXAMPLE.COM" -replacement "\1"
cluster1::> vserver name-mapping insert -vserver vs1 -direction krb-unix -position 3 -pattern
"()(/[^@])?@SEC.EXAMPLE.COM" -replacement "\1 -address 10.238.33.245/24

```
[[ID6f6dfbc11f71a070376efbca52dbf331]]
= vserver name-mapping modify
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a name mapping's pattern, replacement pattern, or both

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver name-mapping modify` command modifies the pattern, the
replacement pattern, or both of a specified name mapping.

You can specify patterns as POSIX regular expressions. For information
about regular expressions, see the UNIX reference page for ``_regex(7)_``
.

Each Vserver can have up to 12500 name mappings in each direction.

NOTE: If you are using the CLI, you must delimit all regular expressions
with double quotation marks ("). For instance, to enter the regular
expression (.+) in the CLI, type "(.+)" at the command prompt. To add a
"?" to the expression, press ESC followed by the "?".



== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to modify the name
mapping.

``-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}`` - Direction::
This parameter specifies the direction of the name mapping. Possible
values are ``_krb-unix_`` for a Kerberos-to-UNIX name mapping, ``_win-
unix_`` for a Windows-to-UNIX name mapping, ``_unix-win_`` for a UNIX-to-
Windows name mapping, ``_s3-unix_`` for a S3-to-UNIX name mapping and
``_s3-win_`` for a S3-to-Windows name mapping.

``-position <integer>`` - Position::
This parameter specifies the name mapping's position in the priority list.
A position is specified as a positive integer. Each mapping direction
(Kerberos-to-UNIX, Windows-to-UNIX, UNIX-to-Windows, S3-to-UNIX and S3-to-
Windows) has its own priority list.

``[-pattern <text>]`` - Pattern::
This parameter specifies the pattern you want to match. Refer to the
```

```
command description section for details. The pattern can be up to 256
characters in length.

``[-replacement <text>]`` - Replacement::
This parameter specifies the replacement pattern. The replacement pattern
can be up to 256 characters in length.

{ ``[-address <IP Address/Mask>]`` - IP Address with Subnet Mask::
This optional parameter specifies the IP address that can be used to match
the client's workstation IP address with the pattern.

| ``[-hostname <text>]`` - Hostname }::
This optional parameter specifies the hostname that can be used to match
the corresponding client's workstation IP address with the list of IP
addresses with the pattern.

== Examples

The following example modifies the name mapping on the Vserver named vs1
and direction win-unix, at position 3. The pattern to be matched is
changed to "EXAMPLE\(.+)".
```

cluster1::> vserver name-mapping modify -vserver vs1 -direction win-unix -position 3 -pattern "EXAMPLE\\(.) -address 10.238.2.54/32" cluster1::> vserver name-mapping modify -vserver vs1 -direction win-unix -position 3 -pattern "EXAMPLE\\(.) -hostname google.com"

```
[[ID26b4854392099f27219cf9739ff915dc]]
= vserver name-mapping refresh-hostname-ip
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Refresh the IP addresses for configured hostnames

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver name-mapping refresh-hostname-ip` command will refresh the IP
Address entries in the name-mapping configuration by resolving the
hostname. If you run this command with no parameters, this will refresh
the IP address entries for every hostname in the name-mapping
configuration.

== Parameters

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the Vserver for which the hostname lookup needs
to be done.

``[-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}]`` - Name
Mapping Direction [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the direction of the name-mapping entry
for the hostname lookup.

``[-hostname <text>]`` - Hostname [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the hostname for which the lookup needs
to be done.

== Examples
```

cluster1::*> vserver name-mapping refresh-hostname-ip -vserver vs1 -direction win-unix -hostname

```
[[ID5ab2ec91b733580c1409b323d434d3a9]]
= vserver name-mapping show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display name mappings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver name-mapping show` command displays information about name
mappings. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information about all name mappings:

* Vserver name
* Direction of the mapping (krb-unix for Kerberos-to-UNIX, win-unix for
Windows-to-UNIX, or unix-win for UNIX-to-Windows)
* Position of the mapping in the priority list
* Pattern to be matched
* Replacement pattern

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
Kerberos-to-UNIX name mappings, run the command with the `-direction krb-
unix` parameter.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the name mapping or mappings that match the specified Vserver.

``[-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}]`` - Direction::
If you specify this parameter, the command displays information only about
the name mapping or mappings that have the specified mapping direction.

``[-position <integer>]`` - Position::
```

```
If you specify this parameter, the command displays information only about
the name mapping that has the specified position in the priority list.

``[-pattern <text>]`` - Pattern::
If you specify this parameter, the command displays information only about
the name mapping or mappings that use the specified matching pattern. The
pattern can be up to 256 characters in length. Refer to the command
description section for details.

``[-replacement <text>]`` - Replacement::
If you specify this parameter, the command displays information only about
the name mapping or mappings that use the specified replacement pattern.

``[-address <IP Address/Mask>]`` - IP Address with Subnet Mask::
If you specify this parameter, the command displays information only about
the name mapping or mappings that use the specified IP address.

``[-hostname <text>]`` - Hostname::
If you specify this parameter, the command displays information only about
the name mapping or mappings that use the specified hostname.

== Examples

The following example displays information about all name mappings:
```

cluster1::> vserver name-mapping show
Vserver: vs1
Direction: win-unix
Position Hostname IP Address/Mask
-------- ---------- ----------
1 google.com - Pattern: EXAMPLE\\administrator
Replacement: nobody
2 - 10.238.2.34/32 Pattern: EXAMPLE\\(.+)
Replacement: _1

```
[[IDc473731d610406d2cbc7b22f5ceb58e5]]
= vserver name-mapping swap
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Exchange the positions of two name mappings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver name-mapping swap` command exchanges the positions of two
name mappings in the priority list.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the name mappings are
located.

``-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}`` - Direction::
This parameter specifies the direction of the name mappings that you want
to exchange. Each mapping direction (Kerberos-to-UNIX, Windows-to-UNIX,
and UNIX-to-Windows) has its own priority list.

``-position <integer>`` - Position::
This parameter specifies the position in the priority list of the first
name mapping that you want to exchange. Specify a position as a positive
integer.

``-with-position <integer>`` - Position of an existing name mapping entry
in the list of name mappings for this Vserver. This entry will be swapped
with the entry at 'position'.::
This parameter specifies the position in the priority list of the second
name mapping that you want to exchange. Specify a position as a positive
integer.

== Examples

The following example exchanges the positions of two name mappings on a
Vserver named vs1. The name mappings have the direction Windows-to-UNIX.
The name mappings are exchanged between positions 2 and 4.
```

cluster1::> vserver name-mapping swap -vserver vs1 -direction win-unix -position 2 -with-position 4

```
:leveloffset: -1


= vserver nfs commands

:leveloffset: +1


[[ID0d8c37311984949abb2ec8c773ba8195]]
= vserver nfs create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an NFS configuration for a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs create` command enables and configures a Vserver to serve
NFS clients. The Vserver must already exist. An NFS-enabled Vserver is
associated with an NIS domain.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to create the NFS
configuration.

``[-access {true|false}]`` - General NFS Access::
This optional parameter specifies whether to enable NFS access on the
Vserver. The default setting is `true` .

``[-rpcsec-ctx-high <integer>]`` - RPC GSS Context Cache High Water Mark
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of RPCSEC_GSS
authentication contexts, which are used by Kerberos. The default setting
is zero. See RFC 2203 for information about RPCSEC_GSS contexts.

``[-rpcsec-ctx-idle <integer>]`` - RPC GSS Context Idle [.advaced-
```

```
priv]#(privilege: advanced)#::
This optional parameter specifies, in seconds, the amount of time a
RPCSEC_GSS context is permitted to remain unused before it is deleted. The
default setting is zero seconds. See RFC 2203 for information about
RPCSEC_GSS contexts.

``[-v3 {enabled|disabled}]`` - NFS v3::
This optional parameter specifies whether to enable access for NFSv3
clients. The default setting is `enabled` .

``[-v4.0 {enabled|disabled}]`` - NFS v4.0::
This optional parameter specifies whether to enable access for NFSv4.0
clients. The default setting is `enabled` .

``[-udp {enabled|disabled}]`` - UDP Protocol::
This optional parameter specifies whether to enable NFS access over UDP.
The default setting is `enabled` .
+
NOTE: Even if UDP is disabled, if TCP is enabled, the Vserver does not
block NFSv3 traffic over UDP. By allowing this traffic, the storage system
can process NFS_NULL ops that the Solaris automounter sends to determine
if the storage system is alive. (Solaris sends these ops over UDP even if
configured to use TCP.) To disallow access for certain clients, including
over UDP, you can use export-policy rules. For more information, see the
xref:{relative_path}vserver-export-policy-rule-create.html[vserver export-
policy rule create] command.
+


``[-tcp {enabled|disabled}]`` - TCP Protocol::
This optional parameter specifies whether to enable NFS access over TCP.
The default setting is `enabled` .

``[-default-win-user <text>]`` - Default Windows User::
This optional parameter specifies a list of default Windows users for the
NFS server.

``[-enable-ejukebox {true|false}]`` - Enable NFSv3 EJUKEBOX error
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether EJUKEBOX errors are enabled for
NFSv3. The default setting is `true` .

``[-v3-require-read-attributes {true|false}]`` - Require All NFSv3 Reads
to Return Read Attributes [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether NFSv3 read operations are
required to return read attributes. The default setting is `false` .
```

```
``[-v3-fsid-change {enabled|disabled}]`` - Show Change in FSID as NFSv3
Clients Traverse Filesystems [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP shows changes in file
system identifiers (FSIDs) as NFSv3 clients traverse file systems. The
default setting is `enabled` .

``[-v3-connection-drop {enabled|disabled}]`` - Enable the Dropping of a
Connection When an NFSv3 Request is Dropped [.advaced-priv]#(privilege:
advanced)#::
This optional parameter specifies whether ONTAP allows to drop the
connection when a NFSv3 request is dropped. The default setting is
`enabled` .

``[-ntfs-unix-security-ops {fail|ignore|use_export_policy}]`` - Vserver
NTFS Unix Security Options [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies how NFSv3 security changes affect NTFS
volumes. If you set this parameter to `ignore` , ONTAP ignores NFSv3
security changes. If you set this parameter to `fail` , this overrides the
unix security options set in the relevant export rules. If you set this
parameter to `use_export_policy` , ONTAP processes NFSv3 security changes
in accordance with the relevant export rules. The default setting is
`use_export_policy` at the time of creation.

``[-chown-mode {restricted|unrestricted|use_export_policy}]`` - Vserver
Change Ownership Mode [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether file ownership can be changed
only by the superuser, or if a non-root user can also change file
ownership. If you set this parameter to `restricted` , file ownership can
be changed only by the superuser, even though the on-disk permissions
allow a non-root user to change file ownership. If you set this parameter
to `unrestricted` , file ownership can be changed by the superuser and by
the non-root user, depending upon the access granted by on-disk
permissions. If you set this parameter to `use_export_policy` , file
ownership can be changed in accordance with the relevant export rules.

``[-trace-enabled {true|false}]`` - NFS Response Trace Enabled [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP logs NFS requests when
they exceed the NFS response trigger time (see the `trigger` parameter).
The default setting is `false` .

``[-trigger <integer>]`` - NFS Response Trigger (in secs) [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the amount of time, in seconds, after
which ONTAP must log an NFS request if it has not completed (assuming the
`-trace-enabled option` is `true` ). The default setting is 60.
```

```
``[-udp-max-xfer-size <integer>]`` - UDP Maximum Transfer Size (bytes)
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum transfer size (in bytes)
that the NFS mount protocol will negotiate with the client for UDP
transport. The range is 8192 to 57344. The default setting is 32768.

``[-tcp-max-xfer-size <integer>]`` - TCP Maximum Transfer Size (bytes)
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum transfer size (in bytes)
that the storage system negotiates with the client for TCP transport of
data for NFSv3, and NFSv4.x protocols. The range is 8192 to 1048576. The
default setting is 65536.
+
NOTE: Setting the parameter value greater than 65536 may cause performance
degradation for existing connections using smaller values. Contact
technical support for guidance.
+



``[-v4.0-acl {enabled|disabled}]`` - NFSv4.0 ACL Support::
This optional parameter specifies whether ONTAP supports NFSv4.0 access
control lists (ACLs). The default setting is `disabled` .

``[-v4.0-read-delegation {enabled|disabled}]`` - NFSv4.0 Read Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4.0 read
delegations. The default setting is `disabled` .

``[-v4.0-write-delegation {enabled|disabled}]`` - NFSv4.0 Write Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4.0 write
delegations. The default setting is `disabled` .

``[-v4-fsid-change {enabled|disabled}]`` - Show Change in FSID as NFSv4
Clients Traverse Filesystems [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP shows changes in file
system identifiers (FSIDs) as NFSv4 clients traverse file systems. The
default setting is `enabled` .
+
NOTE: If users access the storage system using NFSv4 from Solaris 10
clients, you must set this option to `disabled` .
+



``[-v4.0-referrals {enabled|disabled}]`` - NFSv4.0 Referral Support
[.advaced-priv]#(privilege: advanced)#::
```

```
This optional parameter specifies whether ONTAP supports NFSv4.0
referrals. The default setting is `disabled` . You can set this parameter
to enabled only if you also set the `-v4-fsid-change` to `enabled` . If
clients accessing the node do not support NFSv4.0 referrals, set this
option to `disabled` ; otherwise, those clients will not be able to access
the file system.

``[-v4-id-domain <nfs domain>]`` - NFSv4 ID Mapping Domain::
This optional parameter specifies the domain portion of the string form of
user and group names as defined by the NFSv4 protocol. By default, the
domain name is defaultv4iddomain.com. However, the value of this parameter
overrides the default. The domain name must be agreed upon by both the NFS
client and the storage controller before NFSv4 operations can be executed.
It is recommended that the domain be specified in the fully qualified
domain name format.

``[-v4-validate-symlinkdata {enabled|disabled}]`` - NFSv4 Validate UTF-8
Encoding of Symbolic Link Data [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP validates the UTF-8
encoding of symbolic link data. The default setting is `disabled` .

``[-v4-lease-seconds <integer>]`` - NFSv4 Lease Timeout Value (in secs)
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the time period in which ONTAP
irrevocably grants a lock to a client. By default, the lease period is 30
seconds. The minimum value is 10. The maximum value is one less than the
value of the ` -v4-grace-seconds` parameter.

``[-v4-grace-seconds <integer>]`` - NFSv4 Grace Timeout Value (in secs)::
This optional parameter specifies the time period in which clients attempt
to reclaim their locking state from ONTAP during server recovery. By
default, the grace period is 45 seconds. The minimum value is 1 more than
the value of the `-v4-lease-seconds` parameter. The maximum value is 90.

``[-v4-acl-preserve {enabled|disabled}]`` - Preserves and Modifies NFSv4
ACL (and NTFS File Permissions in Unified Security Style)::
This optional parameter specifies if the NFSv4 ACL is preserved or dropped
when chmod is performed. In unified security style, this parameter also
specifies if NTFS file permissions are preserved or dropped when chmod,
chgrp, or chown are performed. The default is `enabled` .

``[-v4.1 {enabled|disabled}]`` - NFSv4.1 Minor Version Support::
This optional parameter specifies whether to enable access for NFSv4.1 or
later clients. The default setting is `enabled` .

``[-rquota {enabled|disabled}]`` - Rquota Enable::
```

```
This optional parameter specifies whether to enable rquota over NFS. The
default setting is `disabled` .

``[-v4.1-implementation-domain <nfs domain>]`` - NFSv4.1 Implementation ID
Domain [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the NFSv4.1 or later implementation
domain.

``[-v4.1-implementation-name <text>]`` - NFSv4.1 Implementation ID Name
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the NFSv4.1 or later implementation
name.

``[-v4.1-implementation-date <Date>]`` - NFSv4.1 Implementation ID Date
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the NFSv4.1 or later implementation
date.

``[-v4.1-pnfs {enabled|disabled}]`` - NFSv4.1 Parallel NFS Support::
This optional parameter specifies whether ONTAP supports parallel NFS over
NFSv4.1 or later. The default setting is `disabled` .

``[-v4.1-referrals {enabled|disabled}]`` - NFSv4.1 Referral Support
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
referrals. The default setting is `disabled` . You can set this parameter
to enabled only if you also set the `-v4-fsid-change` to `enabled` . If
clients accessing the node do not support NFSv4.1 or later referrals, set
this option to `disabled` ; otherwise, those clients will not be able to
access the file system.

``[-v4.1-acl {enabled|disabled}]`` - NFSv4.1 ACL Support::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
access control lists (ACLs). The default setting is `disabled` .

``[-vstorage {enabled|disabled}]`` - NFS vStorage Support::
This optional parameter specifies whether to enable vstorage over NFS. The
default setting is `disabled` .

``[-v4-numeric-ids {enabled|disabled}]`` - NFSv4 Support for Numeric Owner
IDs::
This optional parameter specifies whether the support for numeric string
identifiers in NFSv4 owner attributes is enabled. The default setting is
`enabled` .

``[-default-win-group <text>]`` - Default Windows Group::
This optional parameter specifies a list of default Windows groups for the
```

```
NFS server.

``[-v4.1-read-delegation {enabled|disabled}]`` - NFSv4.1 Read Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
read delegations. The default setting is `disabled` .

``[-v4.1-write-delegation {enabled|disabled}]`` - NFSv4.1 Write Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
write delegations. The default setting is `disabled` .

``[-v4.x-session-num-slots <integer>]`` - Number of Slots in the NFSv4.x
Session slot tables [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the number of entries in the NFSv4.x
session slot table. By default, the number of slots is 180. The maximum
value is 2000.

``[-v4.x-session-slot-reply-cache-size <integer>]`` - Size of the Reply
that will be Cached in Each NFSv4.x Session Slot (in bytes) [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the number of bytes of the reply that
will be cached in each NFSv4.x session slot. By default, the size of the
cached reply is 640 bytes. The maximum value is 4096.

``[-v4-acl-max-aces <integer>]`` - Maximum Number of ACEs per ACL
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of ACEs in an NFSv4
ACL. The range is 192 to 1024. The default value is 400. Setting it to a
value more than the default could cause performance problems for clients
accessing files with NFSv4 ACLs.

``[-mount-rootonly {enabled|disabled}]`` - NFS Mount Root Only::
This optional parameter specifies whether the Vserver allows MOUNT
protocol calls only from privileged ports (port numbers less than 1024).
The default setting is `enabled` .

``[-nfs-rootonly {enabled|disabled}]`` - NFS Root Only::
This optional parameter specifies whether the Vserver allows NFS protocol
calls only from privileged ports (port numbers less than 1024). The
default setting is `disabled` .

``[-auth-sys-extended-groups {enabled|disabled}]`` - AUTH_SYS Extended
Groups Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP supports fetching
auxiliary groups from a name service rather than from the RPC header. The
```

```
default setting is `disabled` .

``[-extended-groups-limit <integer>]`` - AUTH_SYS and RPCSEC_GSS Auxillary
Groups Limit [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of auxiliary groups
supported over RPC security flavors AUTH_SYS and RPCSEC_GSS in ONTAP. The
range is 32 to 1024. The default value is 32.

``[-validate-qtree-export {enabled|disabled}]`` - Validation of Qtree IDs
for Qtree File Operations [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP performs an additional
validation on qtree IDs. The default setting is `enabled` . This parameter
is ignored unless a non-inherited policy has been or is assigned to a
qtree.

``[-mountd-port <integer>]`` - NFS Mount Daemon Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the NFS mount daemon (mountd)
uses. The port numbers allowed are 635 (the default) and 1024 through
9999.

``[-nlm-port <integer>]`` - Network Lock Manager Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the network lock manager
(NLM) uses. The port numbers allowed are 1024 through 9999. The default
setting is `4045` .

``[-nsm-port <integer>]`` - Network Status Monitor Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the network status monitor
(NSM) uses. The port numbers allowed are 1024 through 9999. The default
setting is `4046` .

``[-rquotad-port <integer>]`` - NFS Quota Daemon Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the NFS quota daemon
(rquotad) uses. The port numbers allowed are 1024 through 9999. The
default setting is `4049` .

``[-permitted-enc-types <NFS Kerberos Encryption Type>,...]`` - Permitted
Kerberos Encryption Types::
This optional parameter specifies the permitted encryption types for
Kerberos over NFS. The default setting is `aes-128` ,`aes-256` .

``[-showmount {enabled|disabled}]`` - Showmount Enabled::
 This optional parameter specifies whether to allow or disallow clients to
```

```
see the Vserver's NFS exports list. The default setting is ``_enabled_`` .
+
NOTE: Showmount leverages the MOUNT protocol in NFSv3 to issue an EXPORT
query to the NFS server. If the mount port is not listening or blocked by
a firewall, or if NFSv3 is disabled on the NFS server, showmount queries
fail.
+


``[-name-service-lookup-protocol {TCP|UDP}]`` - Set the Protocol Used for
Name Services Lookups for Exports::
 This optional parameter specifies the protocol to use for doing name
service lookups. The allowed values are `TCP` and `UDP` . The default
setting is `UDP` .

``[-map-unknown-uid-to-default-windows-user {enable|disable}]`` - Map
Unknown UID to Default Windows User [.advaced-priv]#(privilege:
advanced)#::
If you enable this optional parameter, unknown UNIX users that do not have
a name mapping to a Windows user are mapped to the configured default
Windows user. This allows all unknown UNIX users access with the
credentials of the default Windows user. If you disable it, all unknown
UNIX users without name mapping are always denied access. By default, this
parameter is enabled.

``[-netgroup-dns-domain-search {enabled|disabled}]`` - DNS Domain Search
Enabled During Netgroup Lookup [.advaced-priv]#(privilege: advanced)#::
If you enable this optional parameter, during client access check
evaluation in a netgroup, ONTAP performs an additional verification to
ensure that the domain returned from DNS for that client is listed in the
DNS configuration of the Vserver. This enables you to validate the domain
when clients have the same short name in multiple domains. The default
setting is ``_enabled_`` .

``[-netgroup-trust-any-ns-switch-no-match {enabled|disabled}]`` - Trust
No-Match Result from Any Name Service Switch Source During Netgroup Lookup
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies if you can consider a no-match result
from any netgroup ns-switch source to be authoritative. If this option is
enabled, then a no-match response from any one of the netgroup ns-switch
sources is deemed conclusive even if other sources could not be searched.
The default setting is 'disabled', which causes all netgroup ns-switch
sources to be consulted before a no-match result is deemed conclusive.

``[-ntacl-display-permissive-perms {enabled|disabled}]`` - Display maximum
NT ACL Permissions to NFS Client [.advaced-priv]#(privilege: advanced)#::
```

```
This optional parameter controls the permissions that are displayed to
NFSv3 and NFSv4 clients on a file or directory that has an NT ACL set.
When true, the displayed permissions are based on the maximum access
granted by the NT ACL to any user. When false, the displayed permissions
are based on the minimum access granted by the NT ACL to any user. The
default setting is ``_false_`` .

``[-v3-ms-dos-client {enabled|disabled}]`` - NFSv3 MS-DOS Client Support::
This optional parameter specifies whether to enable access for NFSv3 MS-
DOS clients. The default setting is ``_disabled_`` .

``[-ignore-nt-acl-for-root {enabled|disabled}]`` - Ignore the NT ACL Check
for NFS User 'root' [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether Windows ACLs affect root access
from NFS. If this option is enabled, root access from NFS ignores the NT
ACL set on the file or directory. If auditing is enabled for the Vserver
and there is no name-mapping present, then a default SMB credential
(Builtin\administrator) is used for auditing, and an EMS warning is
generated. The default setting is 'disabled', which causes NFS 'root' to
be mapped to a Windows account, like any other NFS user.

``[-cached-cred-positive-ttl <integer>]`` - Time To Live Value (in msecs)
of a Positive Cached Credential [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the age of the positive cached
credentials after which they will be cleared from the cache. The value
specified must be between 60000 and 604800000. The default setting is
`86400000` .

``[-cached-cred-negative-ttl <integer>]`` - Time To Live Value (in msecs)
of a Negative Cached Credential [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the age of the negative cached
credentials after which they will be cleared from the cache. The value
specified must be between 60000 and 604800000. The default setting is
`7200000` .

``[-skip-root-owner-write-perm-check {enabled|disabled}]`` - Skip
Permission Check for NFS Write Calls from Root/Owner [.advaced-
priv]#(privilege: advanced)#::
 This optional parameter specifies if permission checks are to be skipped
for NFS WRITE calls from root/owner. For copying read-only files to a
destination folder which has inheritable ACLs, this option must be
``_enabled_`` . Warning: When ``_enabled_`` , if an NFS client does not
make use of an NFS ACCESS call to check for user-level permissions and
then tries to write onto read-only files, the operation will succeed. The
default setting is ``_disabled_`` .
```

```
``[-v3-64bit-identifiers {enabled|disabled}]`` - Use 64 Bits for NFSv3
FSIDs and File IDs [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP uses 64 bits (instead of
32 bits) for file system identifiers (FSIDs) and file identifiers (file
IDs) that are returned to NFSv3 clients. The default setting is
``_disabled_`` . When `-v3-fsid-change` is disabled, enable this parameter
to avoid file ID collisions.

``[-v4-inherited-acl-preserve {enabled|disabled}]`` - Ignore Client
Specified Mode Bits and Preserve Inherited NFSv4 ACL When Creating New
Files or Directories [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the client-specified mode bits
should be ignored and the inherited NFSv4 ACL should be preserved when
creating new files or directories. The default setting is ``_disabled_`` .

``[-v3-search-unconverted-filename {enabled|disabled}]`` - Fallback to
Unconverted Filename Search [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to continue search without
converting the filename to the Unicode character set while doing lookup in
a directory.

``[-file-session-io-grouping-count <integer>]`` - I/O Count to Be Grouped
as a Session [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the number of read or write operations
on a file from a single client that are grouped and considered as one
session for event generation applications, such as FPolicy. The event is
generated on the first read or write of a file, and subsequently the event
is generated only after the specified `-file-session-io-grouping-count` .
The default value is ``_5000_`` .

``[-file-session-io-grouping-duration <integer>]`` - Duration for I/O to
Be Grouped as a Session (Secs) [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the duration for which the read or write
operationss on a file from a single client are grouped and considered as
one session for event generation applications, such as FPolicy. The
default value is ``_120_`` seconds.

``[-checksum-for-replay-cache {enabled|disabled}]`` - Enable or disable
Checksum for Replay-Cache [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable replay cache checksum
for NFS requests . The default value is ``_enabled_`` .

``[-cached-cred-harvest-timeout <integer>]`` - Harvest timeout (in msecs)
for a Cached Credential [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the harvest timeout for cached
credentials. The value specified must be between 60000 and 604800000. The
```

```
default setting is `86400000` .

``[-idle-connection-timeout <integer>]`` - Idle Connection Timeout Value
(in seconds)::
This optional parameter specifies the idle connection timeout value for
NFS connections in seconds. The value specified must be between 120 and
86400.

``[-allow-idle-connection-timeout {enabled|disabled}]`` - Are Idle NFS
Connections Supported::
This optional parameter specifies whether killing idle NFS connections is
allowed or not. The default setting for `-allow-idle-connection-timeout`
is ``_enabled_`` .

``[-v3-hide-snapshot {enabled|disabled}]`` - Hide Snapshot Directory under
NFSv3 Mount Point::
This optional parameter specifies whether to hide the .snapshot directory
while listing under NFSv3 mount points. However an explicit access to the
.snapshot directory will still be allowed even though the option is
enabled. The default setting is `disabled` .

``[-showmount-rootonly {enabled|disabled}]`` - Provide Root Path as
Showmount State::
This optional parameter specifies whether to provide root path as
showmount state when `-showmount` parameter is disabled. The default value
for `showmount-rootonly` is ``_disabled_`` .

``[-v4-64bit-identifiers {enabled|disabled}]`` - Use 64 Bits for NFSv4.x
FSIDs and File IDs [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP uses 64 bits (instead of
32 bits) for file system identifiers (FSIDs) and file identifiers (file
IDs) that are returned to NFSv4.x clients. The default setting is
``_enabled_`` . When `-v4-fsid-change` is disabled, enable this parameter
to avoid file ID collisions.

``[-v4.2-seclabel {enabled|disabled}]`` - NFSV4.2 Security Label Support
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable security labels for
NFSv4.2. The default setting for `-v4.2-seclabel` is ``_disabled_`` .

``[-rdma {enabled|disabled}]`` - RDMA Protocol::
This optional parameter specifies whether to enable NFS access over RDMA.
The default setting for parameter `-rdma` is ``_enabled_`` .

``[-v4.1-trunking {enabled|disabled}]`` - NFSV4.1 Trunking Support::
This optional parameter specifies whether to enable trunking for NFSv4.1.
```

```
The default setting for `-v4.1-trunking` is ``_disabled_`` .

``[-v4.2-sparsefile-ops {enabled|disabled}]`` - NFSV4.2 Sparse Files Ops
Support [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable sparsefile ops for
NFSv4.2. The default setting for `-v4.2-sparsefile-ops` is ``_disabled_``
.

``[-v4.2-xattrs {enabled|disabled}]`` - NFSV4.2 Extended Attributes
Support::
This optional parameter specifies whether to enable extended attributes
for NFSv4.2. The default setting for `-v4.2-xattrs` is ``_enabled_`` .

``[-v4.1-sequence-admin-revoked {enabled|disabled}]`` - NFSV4 Admin
Revoked Support [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable NFSv4.1 admin revoked
for NFSv4. Once enabled, when any of the client state is marked as revoked
on the server side, the NFS server sends a SEQ4_STATUS_ADMIN_STATE_REVOKED
flag in the SEQUENCE operation reply. The default setting for `-v4.1-
sequence-admin-revoked` is ``_disabled_`` .

``[-v4-acl-compliance {enabled|disabled}]`` - NFSV4 ACL Compliance Support
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable ACL compliance support
for NFSv4 newly created files/directory. Once enabled, the EVERYONE@ ACE
mask will be mapped to UNIX 'owner' and 'group' as well, along with UNIX
'other'. UNIX 'owner' will no longer be considered under GROUP@ ACE. The
default setting for `-v4-acl-compliance` is ``_disabled_`` .

== Examples

The following example enables and configures NFS access on a Vserver named
vs0. NFS access is enabled. The maximum number of RPCSEC_GSS
authentication contexts is set to 5. The RPCSEC_GSS idle time is set to
360 seconds. Access is enabled for NFS v3 clients over both UDP and TCP.
```

```
cluster1::> vserver nfs create -vserver vs0 -access true -rpcsec-ctx-high 5 -rpcsec-ctx-idle 360 -v3 enabled
-udp enabled -tcp enabled
```

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-create.html[vserver
export-policy rule create]




[[ID117e96ba48a0b11bd59f13cd0fed6fc2]]
= vserver nfs delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete the NFS configuration of a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs delete` command deletes the NFS configuration of a
specified Vserver. This command does not delete the Vserver itself, just
its ability to serve NFS clients.

NOTE: If you delete a Vserver, the Vserver's NFS configuration is
automatically deleted. Any Windows-to-UNIX or UNIX-to-Windows name
mappings defined for the Vserver are also deleted because they require
both the CIFS and NFS servers.



== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver whose NFS configuration you want to delete.

== Examples

The following example deletes the NFS configuration of a Vserver named
vs2:
```

cluster1::> vserver nfs delete -vserver vs2

```
[[IDce4572729eb905f0f8ccc62e7a16566c]]
= vserver nfs modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the NFS configuration of a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs modify` command modifies the configuration of an NFS-
enabled Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the Vserver whose NFS configuration you want to modify.

``[-access {true|false}]`` - General NFS Access::
This optional parameter specifies whether NFS access is enabled on the
Vserver.

``[-rpcsec-ctx-high <integer>]`` - RPC GSS Context Cache High Water Mark
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of RPCSEC_GSS
authentication contexts, which are used by Kerberos. The default setting
is zero at the time of creation. See RFC 2203 for information about
RPCSEC_GSS contexts.

``[-rpcsec-ctx-idle <integer>]`` - RPC GSS Context Idle [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies, in seconds, the amount of time a
RPCSEC_GSS context is permitted to remain unused before it is deleted. The
default setting is zero seconds at the time of creation. See RFC 2203 for
information about RPCSEC_GSS contexts.

``[-v3 {enabled|disabled}]`` - NFS v3::
This optional parameter specifies whether to enable access for NFS v3
clients.

``[-v4.0 {enabled|disabled}]`` - NFS v4.0::
This optional parameter specifies whether to enable access for NFSv4.0
```

```
clients. The default setting is `enabled` at the time of creation.

``[-udp {enabled|disabled}]`` - UDP Protocol::
This optional parameter specifies whether to enable NFS access over UDP.
+
NOTE: Even if UDP is disabled, if TCP is enabled, the Vserver does not
block NFSv3 traffic over UDP. By allowing this traffic, the storage system
can process NFS_NULL ops that the Solaris automounter sends to determine
if the storage system is alive. (Solaris sends these ops over UDP even if
configured to use TCP.) To disallow access for certain clients, including
over UDP, you can use export-policy rules. For more information, see the
xref:{relative_path}vserver-export-policy-rule-create.html[vserver export-
policy rule create] command.
+


``[-tcp {enabled|disabled}]`` - TCP Protocol::
This optional parameter specifies whether to enable NFS access over TCP.

``[-default-win-user <text>]`` - Default Windows User::
This optional parameter specifies a list of default Windows users for the
NFS server.

``[-enable-ejukebox {true|false}]`` - Enable NFSv3 EJUKEBOX error
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether EJUKEBOX errors are enabled for
NFSv3. The default setting is `true` at the time of creation.

``[-v3-require-read-attributes {true|false}]`` - Require All NFSv3 Reads
to Return Read Attributes [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether NFSv3 read operations are
required to return read attributes. The default setting is `false` at the
time of creation.

``[-v3-fsid-change {enabled|disabled}]`` - Show Change in FSID as NFSv3
Clients Traverse Filesystems [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP shows changes in file
system identifiers (FSIDs) as NFSv3 clients traverse file systems. If you
change the value of this parameter, clients must remount any paths over
which they are using NFSv3.

``[-v3-connection-drop {enabled|disabled}]`` - Enable the Dropping of a
Connection When an NFSv3 Request is Dropped [.advaced-priv]#(privilege:
advanced)#::
This optional parameter specifies whether NFS v3 connection drop is
enabled. The default setting is `enabled` at the time of creation.
```

```
``[-ntfs-unix-security-ops {fail|ignore|use_export_policy}]`` - Vserver
NTFS Unix Security Options [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies how NFSv3 security changes affect NTFS
volumes. If you set this parameter to `ignore` , ONTAP ignores NFSv3
security changes. If you set this parameter to `fail` , this overrides the
unix security options set in the relevant export rules. If you set this
parameter to `use_export_policy` , ONTAP processes NFSv3 security changes
in accordance with the relevant export rules. The default setting is
`use_export_policy` at the time of creation.

``[-chown-mode {restricted|unrestricted|use_export_policy}]`` - Vserver
Change Ownership Mode [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether file ownership can be changed
only by the superuser, or if a non-root user can also change file
ownership. If you set this parameter to `restricted` , file ownership can
be changed only by the superuser, even though the on-disk permissions
allow a non-root user to change file ownership. If you set this parameter
to `unrestricted` , file ownership can be changed by the superuser and by
the non-root user, depending upon the access granted by on-disk
permissions. If you set this parameter to `use_export_policy` , file
ownership can be changed in accordance with the relevant export rules.

``[-trace-enabled {true|false}]`` - NFS Response Trace Enabled [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP logs NFS requests when
they exceed the NFS response trigger time (see the `trigger` parameter).
The default setting is `false` at the time of creation.

``[-trigger <integer>]`` - NFS Response Trigger (in secs) [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the amount of time, in seconds, after
which ONTAP must log an NFS request if it has not completed (assuming the
`-trace-enabled option` is set to `true` ). The default setting is 60 at
the time of creation.

``[-udp-max-xfer-size <integer>]`` - UDP Maximum Transfer Size (bytes)
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum transfer size (in bytes)
that the NFS mount protocol negotiates with the client for UDP transport.
The range is 8192 to 57344. The default setting is 32768 at the time of
creation.

``[-tcp-max-xfer-size <integer>]`` - TCP Maximum Transfer Size (bytes)
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum transfer size (in bytes)
that the storage system negotiates with the client for TCP transport of
```

```
data for NFSv3 and NFSv4.x protocols. The range is 8192 to 1048576. The
default setting is 65536 when created. Warning: Increasing/decreasing the
value of this parameter could affect the performance for existing
connections.

``[-v4.0-acl {enabled|disabled}]`` - NFSv4.0 ACL Support::
This optional parameter specifies whether ONTAP supports NFSv4.0 access
control lists (ACLs). The default setting is `disabled` when created.

``[-v4.0-read-delegation {enabled|disabled}]`` - NFSv4.0 Read Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4 read
delegations. The default setting is `disabled` when created.

``[-v4.0-write-delegation {enabled|disabled}]`` - NFSv4.0 Write Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4 write
delegations. The default setting is `disabled` when created.

``[-v4-fsid-change {enabled|disabled}]`` - Show Change in FSID as NFSv4
Clients Traverse Filesystems [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP shows changes in file
system identifiers (FSIDs) as NFSv4 clients traverse file systems. The
default setting is `enabled` when created. If you change the value of this
parameter, clients must remount any paths over which they are using NFSv4.
+
NOTE: If users access the storage system using NFSv4 from Solaris 10
clients, you must set this option to `disabled` .
+


``[-v4.0-referrals {enabled|disabled}]`` - NFSv4.0 Referral Support
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP supports NFSv4.0
referrals. The default setting is `disabled` when created. You can set
this parameter to enabled only if the `-v4-fsid-change` option is also set
to `enabled` . If clients accessing the node do not support NFSv4.0
referrals, set this option to `disabled` ; otherwise, those clients will
not be able to access the file system.

``[-v4-id-domain <nfs domain>]`` - NFSv4 ID Mapping Domain::
This optional parameter specifies the domain portion of the string form of
user and group names as defined by the NFSv4 protocol. By default, the
domain name is defaultv4iddomain.com. However, the value of this parameter
overrides the default. The domain name must be agreed upon by both the NFS
client and the storage controller before NFSv4 operations can be executed.
```

```
It is recommended that the domain be specified in the fully qualified
domain name format.

``[-v4-validate-symlinkdata {enabled|disabled}]`` - NFSv4 Validate UTF-8
Encoding of Symbolic Link Data [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP validates the UTF-8
encoding of symbolic link data. The default setting is `disabled` when
created.

``[-v4-lease-seconds <integer>]`` - NFSv4 Lease Timeout Value (in secs)
[.advaced-priv]#(privilege: advanced)#::
This optional parameters specifies the time period in which ONTAP
irrevocably grants a lock to a client. By default, the lease period is 30
seconds. The minimum value is 10. The maximum value is one less than the
value of the ` -v4-grace-seconds` parameter.

``[-v4-grace-seconds <integer>]`` - NFSv4 Grace Timeout Value (in secs)::
This optional parameter specifies the time period in which clients attempt
to reclaim their locking state from ONTAP during server recovery. By
default, the grace period is 45 seconds. The minimum value is 1 more than
the value of the `-v4-lease-seconds` parameter. The maximum value is 90.

``[-v4-acl-preserve {enabled|disabled}]`` - Preserves and Modifies NFSv4
ACL (and NTFS File Permissions in Unified Security Style)::
This optional parameter specifies if the NFSv4 ACL is preserved or dropped
when chmod is performed. In unified security style, this parameter also
specifies if NTFS file permissions are preserved or dropped when chmod,
chgrp, or chown are performed. The default is `enabled` .

``[-v4.1 {enabled|disabled}]`` - NFSv4.1 Minor Version Support::
This optional parameter specifies whether to enable access for NFSv4.1 or
later clients. The default setting is `enabled` at the time of creation.

``[-rquota {enabled|disabled}]`` - Rquota Enable::
This optional parameter specifies whether to enable rquota over NFS. The
default setting is `disabled` at the time of creation.

``[-v4.1-implementation-domain <nfs domain>]`` - NFSv4.1 Implementation ID
Domain [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the NFSv4.1 or later implementation
domain.

``[-v4.1-implementation-name <text>]`` - NFSv4.1 Implementation ID Name
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the NFSv4.1 or later implementation
name.
```

```
``[-v4.1-implementation-date <Date>]`` - NFSv4.1 Implementation ID Date
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the NFSv4.1 or later implementation
date.

``[-v4.1-pnfs {enabled|disabled}]`` - NFSv4.1 Parallel NFS Support::
This optional parameter specifies whether to enable access for pNFS for
NFSv4.1 or later. The default setting is `disabled` at the time of
creation.

``[-v4.1-referrals {enabled|disabled}]`` - NFSv4.1 Referral Support
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
referrals. The default setting is `disabled` when created. You can set
this parameter to enabled only if the `-v4-fsid-change` option is also set
to `enabled` . If clients accessing the node do not support NFSv4.1 or
later referrals, set this option to `disabled` ; otherwise, those clients
will not be able to access the file system.

``[-v4.1-acl {enabled|disabled}]`` - NFSv4.1 ACL Support::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
access control lists (ACLs). The default setting is `disabled` when
created.

``[-vstorage {enabled|disabled}]`` - NFS vStorage Support::
This optional parameter specifies whether to enable vstorage over NFS. The
default setting is `disabled` at the time of creation.

``[-v4-numeric-ids {enabled|disabled}]`` - NFSv4 Support for Numeric Owner
IDs::
This optional parameter specifies whether to enable the support for
numeric string identifiers in NFSv4 owner attributes. The default setting
is `enabled` at the time of creation.

``[-default-win-group <text>]`` - Default Windows Group::
This optional parameter specifies a list of default Windows groups for the
NFS server.

``[-v4.1-read-delegation {enabled|disabled}]`` - NFSv4.1 Read Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
read delegations. The default setting is `disabled` when created.

``[-v4.1-write-delegation {enabled|disabled}]`` - NFSv4.1 Write Delegation
Support::
This optional parameter specifies whether ONTAP supports NFSv4.1 or later
```

```
write delegations. The default setting is `disabled` when created.

``[-v4.x-session-num-slots <integer>]`` - Number of Slots in the NFSv4.x
Session slot tables [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the number of entries in the NFSv4.x
session slot table. By default, the number of slots is 180. The maximum
value is 2000.

``[-v4.x-session-slot-reply-cache-size <integer>]`` - Size of the Reply
that will be Cached in Each NFSv4.x Session Slot (in bytes) [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the number of bytes of the reply that
will be cached in each NFSv4.x session slot. By default, the size of the
cached reply is 640 bytes. The maximum value is 4096.

``[-v4-acl-max-aces <integer>]`` - Maximum Number of ACEs per ACL
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of ACEs in a NFSv4
ACL. The range is 192 to 1024. The default value is 400. Setting it to a
value more than the default could cause performance problems for clients
accessing files with NFSv4 ACLs.

``[-mount-rootonly {enabled|disabled}]`` - NFS Mount Root Only::
This optional parameter specifies whether the Vserver allows MOUNT
protocol calls only from privileged ports (port numbers less than 1024).
The default setting is `enabled` .

``[-nfs-rootonly {enabled|disabled}]`` - NFS Root Only::
This optional parameter specifies whether the Vserver allows NFS protocol
calls only from privileged ports (port numbers less than 1024). The
default setting is `disabled` .

``[-auth-sys-extended-groups {enabled|disabled}]`` - AUTH_SYS Extended
Groups Enabled [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP supports fetching
auxiliary groups from a name service rather than from the RPC header. The
default setting is `disabled` .

``[-extended-groups-limit <integer>]`` - AUTH_SYS and RPCSEC_GSS Auxillary
Groups Limit [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the maximum number of auxiliary groups
supported over RPC security flavors AUTH_SYS and RPCSEC_GSS in ONTAP. The
range is 32 to 1024. The default value is 32.

``[-validate-qtree-export {enabled|disabled}]`` - Validation of Qtree IDs
for Qtree File Operations [.advaced-priv]#(privilege: advanced)#::
```

```
This optional parameter specifies whether ONTAP performs an additional
validation on qtree IDs. The default setting is `enabled` . This parameter
is ignored unless a non-inherited policy has been or is assigned to a
qtree.

``[-mountd-port <integer>]`` - NFS Mount Daemon Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the NFS mount daemon (mountd)
uses. The port numbers allowed are 635 (the default) and 1024 through
9999.

``[-nlm-port <integer>]`` - Network Lock Manager Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the network lock manager
(NLM) uses. The port numbers allowed are 1024 through 9999. The default
setting is `4045` .

``[-nsm-port <integer>]`` - Network Status Monitor Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the network status monitor
(NSM) uses. The port numbers allowed are 1024 through 9999. The default
setting is `4046` .

``[-rquotad-port <integer>]`` - NFS Quota Daemon Port [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies which port the NFS quota daemon
(rquotad) uses. The port numbers allowed are 1024 through 9999. The
default setting is `4049` .

``[-permitted-enc-types <NFS Kerberos Encryption Type>,...]`` - Permitted
Kerberos Encryption Types::
This optional parameter specifies the permitted encryption types for
Kerberos over NFS. The default setting is `aes-128` ,`aes-256` .

``[-showmount {enabled|disabled}]`` - Showmount Enabled::
 This optional parameter specifies whether to allow or disallow clients to
see the Vserver's NFS exports list. The default setting is ``_enabled_`` .
+
NOTE: Showmount leverages the MOUNT protocol in NFSv3 to issue an EXPORT
query to the NFS server. If the mount port is not listening or blocked by
a firewall, or if NFSv3 is disabled on the NFS server, showmount queries
fail.
+


``[-name-service-lookup-protocol {TCP|UDP}]`` - Set the Protocol Used for
```

```
Name Services Lookups for Exports::
 This optional parameter specifies the protocol to use for doing name
service lookups. The allowed values are `TCP` and `UDP` . The default
setting is `UDP` .

``[-map-unknown-uid-to-default-windows-user {enable|disable}]`` - Map
Unknown UID to Default Windows User [.advaced-priv]#(privilege:
advanced)#::
If you enable this optional parameter, unknown UNIX users that do not have
a name mapping to a Windows user are mapped to the configured default
Windows user. This allows all unknown UNIX users access with the
credentials of the default Windows user. If you disable it, all unknown
UNIX users without name mapping are always denied access. By default, this
parameter is enabled.

``[-netgroup-dns-domain-search {enabled|disabled}]`` - DNS Domain Search
Enabled During Netgroup Lookup [.advaced-priv]#(privilege: advanced)#::
If you enable this optional parameter, during client access check
evaluation in a netgroup, ONTAP performs an additional verification to
ensure that the domain returned from DNS for that client is listed in the
DNS configuration of the Vserver. This enables you to validate the domain
when clients have the same short name in multiple domains. The default
setting is ``_enabled_`` .

``[-netgroup-trust-any-ns-switch-no-match {enabled|disabled}]`` - Trust
No-Match Result from Any Name Service Switch Source During Netgroup Lookup
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies if you can consider a no-match result
from any of the netgroup ns-switch sources to be authoritative. If this
option is enabled, then a no-match response from any of the netgroup ns-
switch sources is deemed conclusive even if other sources could not be
searched. The default setting is 'disabled', which causes all netgroup ns-
switch sources to be consulted before a no-match result is deemed
conclusive.

``[-ntacl-display-permissive-perms {enabled|disabled}]`` - Display maximum
NT ACL Permissions to NFS Client [.advaced-priv]#(privilege: advanced)#::
This optional parameter controls the permissions that are displayed to
NFSv3 and NFSv4 clients on a file or directory that has an NT ACL set.
When true, the displayed permissions are based on the maximum access
granted by the NT ACL to any user. When false, the displayed permissions
are based on the minimum access granted by the NT ACL to any user. The
default setting is ``_false_`` .

``[-v3-ms-dos-client {enabled|disabled}]`` - NFSv3 MS-DOS Client Support::
This optional parameter specifies whether to enable access for NFSv3 MS-
```

```
DOS clients. The default setting is ``_disabled_`` at the time of
creation.

``[-ignore-nt-acl-for-root {enabled|disabled}]`` - Ignore the NT ACL Check
for NFS User 'root' [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether Windows ACLs affect root access
from NFS. If this option is enabled, root access from NFS ignores the NT
ACL set on the file or directory. If auditing is enabled for the Vserver
and there is no name-mapping present, then a default SMB credential
(Builtin\administrator) is used for auditing, and an EMS warning is
generated. The default setting is 'disabled', which causes NFS 'root' to
be mapped to a Windows account, like any other NFS user.

``[-cached-cred-positive-ttl <integer>]`` - Time To Live Value (in msecs)
of a Positive Cached Credential [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the age of the positive cached
credentials after which they will be cleared from the cache. The value
specified must be between 60000 and 604800000. The default setting is
`86400000` .

``[-cached-cred-negative-ttl <integer>]`` - Time To Live Value (in msecs)
of a Negative Cached Credential [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the age of the negative cached
credentials after which they will be cleared from the cache. The value
specified must be between 60000 and 604800000. The default setting is
`7200000` .

``[-skip-root-owner-write-perm-check {enabled|disabled}]`` - Skip
Permission Check for NFS Write Calls from Root/Owner [.advaced-
priv]#(privilege: advanced)#::
 This optional parameter specifies if permission checks are to be skipped
for NFS WRITE calls from root/owner. For copying read-only files to a
destination folder which has inheritable ACLs, this option must be
``_enabled_`` . Warning: When ``_enabled_`` , if an NFS client does not
make use of an NFS ACCESS call to check for user-level permissions and
then tries to write onto read-only files, the operation will succeed. The
default setting is ``_disabled_`` .

``[-v3-64bit-identifiers {enabled|disabled}]`` - Use 64 Bits for NFSv3
FSIDs and File IDs [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP uses 64 bits (instead of
32 bits) for file system identifiers (FSIDs) and file identifiers (file
IDs) that are returned to NFSv3 clients. If you change the value of this
parameter, clients must remount any paths over which they are using NFSv3.
When `-v3-fsid-change` is disabled, enable this parameter to avoid file ID
collisions.
```

```
``[-v4-inherited-acl-preserve {enabled|disabled}]`` - Ignore Client
Specified Mode Bits and Preserve Inherited NFSv4 ACL When Creating New
Files or Directories [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether the client-specified mode bits
should be ignored and the inherited NFSv4 ACL should be preserved when
creating new files or directories. The default setting is ``_disabled_`` .

``[-v3-search-unconverted-filename {enabled|disabled}]`` - Fallback to
Unconverted Filename Search [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to continue search without
converting the filename to the Unicode character set while doing lookup in
a directory.

``[-file-session-io-grouping-count <integer>]`` - I/O Count to Be Grouped
as a Session [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the number of read or write operations
on a file from a single client that are grouped and considered as one
session for event generation applications, such as FPolicy. The event is
generated on the first read or write of a file, and subsequently the event
is generated only after the specified `-file-session-io-grouping-count` .
The default value is ``_5000_`` .

``[-file-session-io-grouping-duration <integer>]`` - Duration for I/O to
Be Grouped as a Session (Secs) [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the duration for which the read or write
operationss on a file from a single client are grouped and considered as
one session for event generation applications, such as FPolicy. The
default value is ``_120_`` seconds.

``[-checksum-for-replay-cache {enabled|disabled}]`` - Enable or disable
Checksum for Replay-Cache [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable replay cache checksum
for NFS requests. The default value is ``_enabled_`` .

``[-cached-cred-harvest-timeout <integer>]`` - Harvest timeout (in msecs)
for a Cached Credential [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the harvest timeout for cached
credentials. The value specified must be between 60000 and 604800000. The
default setting is `86400000` .

``[-idle-connection-timeout <integer>]`` - Idle Connection Timeout Value
(in seconds)::
This optional parameter specifies the idle connection timeout for NFS
connections. The value specified must be between 120 and 86400.

``[-allow-idle-connection-timeout {enabled|disabled}]`` - Are Idle NFS
```

```
Connections Supported::
This optional parameter specifies whether killing idle NFS connections is
allowed or not. The default setting for `-allow-idle-connection-timeout`
is ``_enabled_`` .

``[-v3-hide-snapshot {enabled|disabled}]`` - Hide Snapshot Directory under
NFSv3 Mount Point::
This optional parameter specifies whether to hide the .snapshot directory
while listing under NFSv3 mount points. However an explicit access to the
.snapshot directory will still be allowed even though the option is
enabled. The default setting is `disabled` at the time of creation.

``[-showmount-rootonly {enabled|disabled}]`` - Provide Root Path as
Showmount State::
This optional parameter specifies whether to provide root path as
showmount state when `-showmount` parameter is disabled. The default value
for `showmount-rootonly` is ``_disabled_`` .

``[-v4-64bit-identifiers {enabled|disabled}]`` - Use 64 Bits for NFSv4.x
FSIDs and File IDs [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether ONTAP uses 64 bits (instead of
32 bits) for file system identifiers (FSIDs) and file identifiers (file
IDs) that are returned to NFSv4.x clients. If you change the value of this
parameter, clients must remount any paths over which they are using
NFSv4.x. When `-v4-fsid-change` is disabled, enable this parameter to
avoid file ID collisions.

``[-v4.2-seclabel {enabled|disabled}]`` - NFSV4.2 Security Label Support
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable security labels for
NFSv4.2. The default setting for `-v4.2-seclabel` is ``_disabled_`` .

``[-rdma {enabled|disabled}]`` - RDMA Protocol::
This optional parameter specifies whether to enable NFS access over RDMA.
The default setting for parameter `-rdma` is ``_enabled_`` .

``[-v4.1-trunking {enabled|disabled}]`` - NFSV4.1 Trunking Support::
This optional parameter specifies whether to enable trunking for NFSv4.1.
The default setting for `-v4.1-trunking` is ``_disabled_`` .
+
NOTE: If the `-v4.1-trunking` is set to ``_enabled_`` and if users access
the storage system using NFSv4.1 with NFS clients that don't support
session trunking, they may see performance drop, due to use of a single
TCP connection for the multiple mounts to the SVM data LIFs.
+
```

```
+
NOTE: If the `-v4.1-trunking` is changed from ``_enabled_`` to
``_disabled_`` and if users access the storage system using NFSv4 with
trunking mount then client may see IO hung issue, because on the client
side session trunking is still in use and on the Nfs4 server side it is
getting disabled. Remount is required on the Nfs client if such issue
happens.
+


``[-v4.2-sparsefile-ops {enabled|disabled}]`` - NFSV4.2 Sparse Files Ops
Support [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable sparsefile ops for
NFSv4.2. The default setting for `-v4.2-sparsefile-ops` is ``_disabled_``
.

``[-v4.2-xattrs {enabled|disabled}]`` - NFSV4.2 Extended Attributes
Support::
This optional parameter specifies whether to enable extended attributes
for NFSv4.2. The default setting for `-v4.2-xattrs` is ``_enabled_`` .

``[-v4.1-sequence-admin-revoked {enabled|disabled}]`` - NFSV4 Admin
Revoked Support [.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable admin revoked sequence
flag support for NFSv4. Once enabled, when any of the client state is
marked as revoked on the server side, the NFS server sends a
SEQ4_STATUS_ADMIN_STATE_REVOKED flag in the SEQUENCE operation reply. The
default setting for `-v4.1-sequence-admin-revoked` is ``_disabled_`` .

``[-v4-acl-compliance {enabled|disabled}]`` - NFSV4 ACL Compliance Support
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies whether to enable ACL compliance support
for NFSv4 newly created files/directory. Once enabled, the EVERYONE@ ACE
mask will be mapped to UNIX 'owner' and 'group' as well, along with UNIX
'other'. UNIX 'owner' will no longer be considered under GROUP@ ACE. The
default setting for `-v4-acl-compliance` is ``_disabled_`` .

== Examples

The following example enables NFS access on a Vserver named vs0 for NFS
clients that use NFS v3 over TCP:
```

cluster1::> vserver nfs modify -vserver vs0 -access true -v3 enabled -udp disabled -tcp enabled

```
== Related Links

* xref:{relative_path}vserver-export-policy-rule-create.html[vserver
export-policy rule create]




[[IDeaebb0e4bc78058f8c9e552f25c8ce85]]
= vserver nfs off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable the NFS service of a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs off` command disables NFS access on a Vserver. The
Vserver must already exist.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to disable NFS
access.

== Examples

The following example disables NFS access on a Vserver named vs0.
```

cluster1::> vserver nfs off -vserver vs0

```
[[ID97e5868b35039fd761377179c1d8be6e]]
= vserver nfs on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable the NFS service of a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs on` command enables NFS access on a Vserver. The Vserver
must already exist.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to enable NFS
access.

== Examples

The following example enables NFS access on a Vserver named vs0.
```

cluster1::> vserver nfs on -vserver vs0

```
[[ID180ba86167e6263de543de279906b430]]
= vserver nfs prepare-for-v3-ms-dos-client-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Disable NFSv3 MS-DOS Client Support

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description


NOTE: This command is deprecated and may be removed in a future release.

The `vserver nfs prepare-for-v3-ms-dos-client-downgrade` command verifies
that the NFSv3 MS-DOS client setting is disabled on all Vservers and
disables the NFSv3 MS-DOS client support capability on the cluster when
downgrading ONTAP to a version that does not support NFSv3 MS-DOS clients.

== Examples

The following example disables NFSv3 MS-DOS client support on the
Vservers.
```

cluster::1> vserver nfs prepare-for-v3-ms-dos-client-downgrade

```
[[ID73ecf81c946df82f057915e9b078fc8c]]
= vserver nfs prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove NFS configurations that are not compatible with earlier versions of
Data ONTAP

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver nfs prepare-to-downgrade` command removes NFS configurations
incompatible with the earlier release of ONTAP.

== Parameters

``-disable-feature-set <downgrade version>`` - Data ONTAP Version
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the ONTAP version that introduced the new NFS
configurations and needs to be removed before downgrade. The value can be
one of the following:
+
  * 9.2.0 - Remove the NFS configurations introduced in ONTAP release
9.2.0. The configurations include the following:
+
    * `-file-session-io-grouping-count` .
    * `-file-session-io-grouping duration` .





== Examples
```

cluster1::*> vserver nfs prepare-to-downgrade -disable-feature-set 9.2.0

```
[[ID674491e2a761590b7b33d551114f583f]]
= vserver nfs show
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the NFS configurations of Vservers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs show` command displays information about NFS-enabled
Vservers. The command output depends on the parameter or parameters
specified with the command. If no parameters are specified, the command
displays the following information about all NFS-enabled Vservers:

* Vserver name
* Whether general NFS access is enabled
* Whether access to NFSv3 clients is enabled
* Whether access to NFSv4 clients is enabled
* Whether NFS access over UDP is enabled
* Whether NFS access over TCP is enabled
* List of default Windows users (detailed view only)

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
Vservers that enable access over TCP, enter the command with the `-tcp-
enable true` parameter. + + All NFSv4.1 options apply to NFSv4.1 or later
versions.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-krb-opts  ]`` [.advaced-priv]#(privilege: advanced)#::
If you specify the parameter for `-instance` , the command shows detailed
information about all NFS-enabled Vservers. Otherwise, if the `-krb-opts`
parameter is specified, the command shows the following Kerberos-related
information:
+
  * Vserver name
  * Maximum number of RPCSEC_GSS authentication contexts
  * Time, in seconds, an RPCSEC_GSS context can remain idle before being
deleted
```

```
+
Otherwise, if the `-fields` parameter is specified, the command shows
information about all of the NFS-enabled Vservers that you specify as a
comma-delimited list.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the specified NFS-enabled Vserver.

``[-access {true|false}]`` - General NFS Access::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers that have the specified general-access setting.

``[-rpcsec-ctx-high <integer>]`` - RPC GSS Context Cache High Water Mark
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
NFS-enabled Vservers that have the specified maximum number of RPCSEC_GSS
authentication contexts.

``[-rpcsec-ctx-idle <integer>]`` - RPC GSS Context Idle [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers that have the specified timeout value for idle
RPCSEC_GSS contexts.

``[-v3 {enabled|disabled}]`` - NFS v3::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v3` option matches the specified
input.

``[-v4.0 {enabled|disabled}]`` - NFS v4.0::
If you specify this parameter, the command displays information only about
NFS-enabled Vservers for which the `v4.0` option matches the specified
input.

``[-udp {enabled|disabled}]`` - UDP Protocol::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers that have the specified NFS-over-UDP access
setting.

``[-tcp {enabled|disabled}]`` - TCP Protocol::
If you specify this parameter, the command displays information only about
```

```
the NFS-enabled Vservers that have the specified NFS-over-TCP setting.

``[-default-win-user <text>]`` - Default Windows User::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers that have the specified list of default Windows
users.

``[-enable-ejukebox {true|false}]`` - Enable NFSv3 EJUKEBOX error
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `enable-ejukebox` option matches
the specified input.

``[-v3-require-read-attributes {true|false}]`` - Require All NFSv3 Reads
to Return Read Attributes [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which NFSv3 read operations are required or
not required to return read attributes.

``[-v3-fsid-change {enabled|disabled}]`` - Show Change in FSID as NFSv3
Clients Traverse Filesystems [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about
changes in file system identifiers (FSIDs) as NFSv3 clients traverse file
systems.

``[-v3-connection-drop {enabled|disabled}]`` - Enable the Dropping of a
Connection When an NFSv3 Request is Dropped [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v3-connection-drop` option matches
the specified input.

``[-ntfs-unix-security-ops {fail|ignore|use_export_policy}]`` - Vserver
NTFS Unix Security Options [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the NTFS-UNIX security setting matches
the specified input.

``[-chown-mode {restricted|unrestricted|use_export_policy}]`` - Vserver
Change Ownership Mode [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the chown-mode setting matches the
specified input.

``[-trace-enabled {true|false}]`` - NFS Response Trace Enabled [.advaced-
priv]#(privilege: advanced)#::
```

```
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `trace-enabled` option matches the
specified input.

``[-trigger <integer>]`` - NFS Response Trigger (in secs) [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers with the specified NFS response trigger time.

``[-udp-max-xfer-size <integer>]`` - UDP Maximum Transfer Size (bytes)
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers with the specified UDP maximum transfer size. The
range is 8192 to 57344.

``[-tcp-max-xfer-size <integer>]`` - TCP Maximum Transfer Size (bytes)
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers with the specified TCP maximum transfer size. The
range is 8192 to 1048576.

``[-v4.0-acl {enabled|disabled}]`` - NFSv4.0 ACL Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.0-acl` option matches the
specified input.

``[-v4.0-read-delegation {enabled|disabled}]`` - NFSv4.0 Read Delegation
Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.0-read-delegation` option
matches the specified input.

``[-v4.0-write-delegation {enabled|disabled}]`` - NFSv4.0 Write Delegation
Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.0-write-delegation` option
matches the specified input.

``[-v4-fsid-change {enabled|disabled}]`` - Show Change in FSID as NFSv4
Clients Traverse Filesystems [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the showing of NFSv4 file system
identifier (FSID) changes has been enabled or disabled.

``[-v4.0-referrals {enabled|disabled}]`` - NFSv4.0 Referral Support
[.advaced-priv]#(privilege: advanced)#::
```

```
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.0-referrals` option matches the
specified input.

``[-v4-id-domain <nfs domain>]`` - NFSv4 ID Mapping Domain::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers having the specified domain name.

``[-v4-validate-symlinkdata {enabled|disabled}]`` - NFSv4 Validate UTF-8
Encoding of Symbolic Link Data [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which validation of UTF-8 encoding of
symbolic link data has been enabled or disabled.

``[-v4-lease-seconds <integer>]`` - NFSv4 Lease Timeout Value (in secs)
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, it displays the locking lease period. It is
expressed in seconds. Clients that have been inactive for a period equal
or longer to the lease period may lose all their locking state on a node.

``[-v4-grace-seconds <integer>]`` - NFSv4 Grace Timeout Value (in secs)::
If you specify this parameter, it displays the grace period for clients to
reclaim file locks after a server failure. The grace period is expressed
in seconds.

``[-v4-acl-preserve {enabled|disabled}]`` - Preserves and Modifies NFSv4
ACL (and NTFS File Permissions in Unified Security Style)::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4-acl-preserve` option matches
the specified input.

``[-v4.1 {enabled|disabled}]`` - NFSv4.1 Minor Version Support::
If you specify this parameter, the command displays information only about
NFS-enabled Vservers for which the `v4.1` option matches the specified
input.

``[-rquota {enabled|disabled}]`` - Rquota Enable::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `rquota` option matches the
specified input.

``[-v4.1-implementation-domain <nfs domain>]`` - NFSv4.1 Implementation ID
Domain [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-implementation-domain` option
matches the specified input.
```

```
``[-v4.1-implementation-name <text>]`` - NFSv4.1 Implementation ID Name
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-implementation-name` option
matches the specified input.

``[-v4.1-implementation-date <Date>]`` - NFSv4.1 Implementation ID Date
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-implementation-date` option
matches the specified input.

``[-v4.1-pnfs {enabled|disabled}]`` - NFSv4.1 Parallel NFS Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-pnfs` option matches the
specified input.

``[-v4.1-referrals {enabled|disabled}]`` - NFSv4.1 Referral Support
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-referrals` option matches the
specified input.

``[-v4.1-acl {enabled|disabled}]`` - NFSv4.1 ACL Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-acl` option matches the
specified input.

``[-vstorage {enabled|disabled}]`` - NFS vStorage Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `vstorage` option matches the
specified input.

``[-v4-numeric-ids {enabled|disabled}]`` - NFSv4 Support for Numeric Owner
IDs::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4-numeric-ids` option matches the
specified input.

``[-default-win-group <text>]`` - Default Windows Group::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers that have the specified list of default Windows
groups.

``[-v4.1-read-delegation {enabled|disabled}]`` - NFSv4.1 Read Delegation
Support::
```

```
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-read-delegation` option
matches the specified input.

``[-v4.1-write-delegation {enabled|disabled}]`` - NFSv4.1 Write Delegation
Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-write-delegation` option
matches the specified input.

``[-v4.x-session-num-slots <integer>]`` - Number of Slots in the NFSv4.x
Session slot tables [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, this command displays information only
about the NFS-enabled Vservers for which the `v4.x-session-num-slots`
option matches the specified input. The range is 1 to 2000.

``[-v4.x-session-slot-reply-cache-size <integer>]`` - Size of the Reply
that will be Cached in Each NFSv4.x Session Slot (in bytes) [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, this command displays information only
about the NFS-enabled Vservers for which the `v4.x-session-slot-reply-
cache-size` option matches the specified input. The cache size is
expressed in bytes. The range is 512 to 4096.

``[-v4-acl-max-aces <integer>]`` - Maximum Number of ACEs per ACL
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4-acl-max-aces` option matches
the specified input.

``[-mount-rootonly {enabled|disabled}]`` - NFS Mount Root Only::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `mount-rootonly` option matches the
specified input.

``[-nfs-rootonly {enabled|disabled}]`` - NFS Root Only::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `nfs-rootonly` option matches the
specified input.

``[-auth-sys-extended-groups {enabled|disabled}]`` - AUTH_SYS Extended
Groups Enabled [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `auth-sys-extended-groups` option
matches the specified input.
```

```
``[-extended-groups-limit <integer>]`` - AUTH_SYS and RPCSEC_GSS Auxillary
Groups Limit [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers for which the `extended-groups-limit` option matches
the specified input. The range is 32 to 1024.

``[-validate-qtree-export {enabled|disabled}]`` - Validation of Qtree IDs
for Qtree File Operations [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `validate-qtree-export` option matches
the specified input.

``[-mountd-port <integer>]`` - NFS Mount Daemon Port [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `mountd-port` option matches the
specified input.

``[-nlm-port <integer>]`` - Network Lock Manager Port [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `nlm-port` option matches the
specified input.

``[-nsm-port <integer>]`` - Network Status Monitor Port [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `nsm-port` option matches the
specified input.

``[-rquotad-port <integer>]`` - NFS Quota Daemon Port [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `rquotad-port` option matches the
specified input.

``[-permitted-enc-types <NFS Kerberos Encryption Type>,...]`` - Permitted
Kerberos Encryption Types::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `permitted-enc-types` option matches
any of the following : aes-128, aes-256.

``[-showmount {enabled|disabled}]`` - Showmount Enabled::
 If you specify this parameter, the command displays information only
about the NFS-enabled Vserver's for which the `showmount` option matches
the specified input.
```

```
``[-name-service-lookup-protocol {TCP|UDP}]`` - Set the Protocol Used for
Name Services Lookups for Exports::
 If you specify this parameter, the command displays information only
about the NFS-enabled Vservers for which `-name-service-lookup-protocol`
matches the parameter.

``[-map-unknown-uid-to-default-windows-user {enable|disable}]`` - Map
Unknown UID to Default Windows User [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `-map-unknown-uid-to-default-windows-
user` is enabled or disabled.

``[-netgroup-dns-domain-search {enabled|disabled}]`` - DNS Domain Search
Enabled During Netgroup Lookup [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `-netgroup-dns-domain-search` is
enabled or disabled.

``[-netgroup-trust-any-ns-switch-no-match {enabled|disabled}]`` - Trust
No-Match Result from Any Name Service Switch Source During Netgroup Lookup
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `-netgroup-trust-any-ns-switch-no-
match` is enabled or disabled.

``[-ntacl-display-permissive-perms {enabled|disabled}]`` - Display maximum
NT ACL Permissions to NFS Client [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `-ntacl-display-permissive-perms`
matches the parameter.

``[-v3-ms-dos-client {enabled|disabled}]`` - NFSv3 MS-DOS Client Support::
If you specify this parameter, the command displays information only about
NFS-enabled Vservers for which NFSv3 MS-DOS client support is enabled or
disabled.

``[-ignore-nt-acl-for-root {enabled|disabled}]`` - Ignore the NT ACL Check
for NFS User 'root' [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `-ignore-nt-acl-for-root` is enabled or
disabled.

``[-cached-cred-positive-ttl <integer>]`` - Time To Live Value (in msecs)
of a Positive Cached Credential [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
```

```
NFS-enabled Vservers time to live value of the positive cached
credentials.

``[-cached-cred-negative-ttl <integer>]`` - Time To Live Value (in msecs)
of a Negative Cached Credential [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers time to live value of the negative cached
credentials.

``[-skip-root-owner-write-perm-check {enabled|disabled}]`` - Skip
Permission Check for NFS Write Calls from Root/Owner [.advaced-
priv]#(privilege: advanced)#::
 If you specify this parameter, the command displays information only
about the NFS-enabled Vservers for which `-skip-root-owner-write-perm-
check` is enabled or disabled.

``[-v3-64bit-identifiers {enabled|disabled}]`` - Use 64 Bits for NFSv3
FSIDs and File IDs [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `-v3-64bit-identifiers` is enabled or
disabled.

``[-v4-inherited-acl-preserve {enabled|disabled}]`` - Ignore Client
Specified Mode Bits and Preserve Inherited NFSv4 ACL When Creating New
Files or Directories [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers for which `-v4-inherited-acl-preserve` matches the
specified input.

``[-v3-search-unconverted-filename {enabled|disabled}]`` - Fallback to
Unconverted Filename Search [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers for which `-v3-search-unconverted-filename` matches
the specified input.

``[-file-session-io-grouping-count <integer>]`` - I/O Count to Be Grouped
as a Session [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled SVMs for which the `-file-session-io-grouping-count` matches
the specified input.

``[-file-session-io-grouping-duration <integer>]`` - Duration for I/O to
Be Grouped as a Session (Secs) [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled SVMs for which the `-file-session-io-grouping-duration`
matches the specified input.
```

```
``[-checksum-for-replay-cache {enabled|disabled}]`` - Enable or disable
Checksum for Replay-Cache [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled SVMs for which the `-checksum-for-replay-cache` matches the
specified input.

``[-cached-cred-harvest-timeout <integer>]`` - Harvest timeout (in msecs)
for a Cached Credential [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers harvest timeout for cached credentials.

``[-idle-connection-timeout <integer>]`` - Idle Connection Timeout Value
(in seconds)::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers idle connections timeout

``[-allow-idle-connection-timeout {enabled|disabled}]`` - Are Idle NFS
Connections Supported::
If you specify this parameter, the command displays information only about
NFS-enabled Vservers for which the `-allow-idle-connection-timeout` option
matches the specified input.

``[-v3-hide-snapshot {enabled|disabled}]`` - Hide Snapshot Directory under
NFSv3 Mount Point::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers for which `-v3-hide-snapshot` matches the specified
input.

``[-showmount-rootonly {enabled|disabled}]`` - Provide Root Path as
Showmount State::
If you specify this parameter, the command displays information about the
NFS-enabled Vservers for which `-showmount-rootonly` matches the specified
input.

``[-v4-64bit-identifiers {enabled|disabled}]`` - Use 64 Bits for NFSv4.x
FSIDs and File IDs [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which `-v4-64bit-identifiers` is enabled or
disabled.

``[-v4.2-seclabel {enabled|disabled}]`` - NFSV4.2 Security Label Support
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.2-seclabel` option matches the
specified input.
```

```
``[-rdma {enabled|disabled}]`` - RDMA Protocol::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers that have the specified NFS-over-RDMA setting.

``[-v4.1-trunking {enabled|disabled}]`` - NFSV4.1 Trunking Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.1-trunking` option matches the
specified input.

``[-v4.2-sparsefile-ops {enabled|disabled}]`` - NFSV4.2 Sparse Files Ops
Support [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.2-sparsefile-ops` option
matches the specified input.

``[-v4.2-xattrs {enabled|disabled}]`` - NFSV4.2 Extended Attributes
Support::
If you specify this parameter, the command displays information only about
the NFS-enabled Vservers for which the `v4.2-xattrs` option matches the
specified input.

``[-v4.1-sequence-admin-revoked {enabled|disabled}]`` - NFSV4 Admin
Revoked Support [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command only displays information about
the NFS-enabled Vservers for which the `v4.1-sequence-admin-revoked`
option matches the specified input.

``[-v4-acl-compliance {enabled|disabled}]`` - NFSV4 ACL Compliance Support
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command only displays information about
the NFS-enabled Vservers where NFSv41 acl compliance is enabled or
disabled.

== Examples

The following example displays information about all NFS-enabled Vservers:
```

cluster1::> vserver nfs show
General Default
Vserver Access v3 v4 v4.1 UDP TCP Windows User
---------- ------- -------- -------- -------- -------- -------- ----------
vs0 true enabled disabled disabled enabled enabled -
vs1 true enabled disabled disabled enabled enabled -
2 entries were displayed.

```
The following example displays Kerberos-related information about all NFS-
enabled Vservers:
```

cluster1::*> vserver nfs show -krb-opts
Vserver Context High Context Idle
--- --- ---
vs0 30 30
vs1 30 30
2 entries were displayed.

```
[[ID92bd6154888552f6749e6d0c8e4b9568]]
= vserver nfs start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start the NFS service of a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs start` command starts the NFS service on a Vserver to
serve NFS clients. The Vserver must already exist.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to start the NFS
service.

== Examples

The following example starts the NFS service on a Vserver named vs0.
```

cluster1::> vserver nfs start -vserver vs0

```
[[IDae44ec2bbf0f5503c9d18110beaae834]]
= vserver nfs status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the status of the NFS service of a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs status` command shows the status of NFS on a Vserver. The
Vserver must already exist.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver for which you want to see the NFS
status.

``[-is-enabled {true|false}]`` - NFS Service Enabled::
If you specify this optional parameter, the command displays whether NFS
is enabled or not. This parameter is true if the NFS server is running.

== Examples

The following example shows the status of NFS on a Vserver named vs0 for
which NFS is enabled.
```

cluster1::> vserver nfs status -vserver vs0.
The NFS server is running.

```
[[IDdec559d56f7ff8c053bd2389524c7f76]]
= vserver nfs stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop the NFS service of a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs stop` command stops the NFS service on a Vserver to serve
NFS clients. The Vserver must already exist.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to stop the NFS
service.

== Examples

The following example stops the NFS service on a Vserver named vs0.
```

cluster1::> vserver nfs stop -vserver vs0

```
[[IDa3667ac1353e47f275022b51b1cbb75f]]
= vserver nfs connected-clients show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display List of recent NFS clients information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `nfs connected-clients show ` command displays information about
currently connected NFS clients, other NFS clients that are are currently
idle but can be connected, and a list of recently unmounted clients. If a
client connected to the NFS server is idle for longer than the maximum
cache idle time, then the entry will be removed. By default, the maximum
cache idle time is 48 hours.

== 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>|local}]`` - Node Name::
Use this parameter to display information only about the specified node.

``[-vserver <vserver>]`` - Vserver::
Use this parameter to display information only about the specified
vserver.

``[-data-lif-ip <IP Address>]`` - Data LIF IP Address::
Use this parameter to display information only about the specified data
LIF IP.

``[-client-ip <IP Address>]`` - Client IP Address::
Use this parameter to display information only about the specified client
IP.

``[-volume <volume name>]`` - Volume Accessed::
Use this parameter to display information only about the specified volume.

``[-protocol <Client Access Protocol>]`` - Protocol Version::
Use this parameter to display information only about the specified
protocol.

``[-idle-time <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Idle
Time (Sec)::
Use this parameter to display the time elapsed since the client sent the
last request for this volume.

``[-local-reqs <integer>]`` - Number of Local Reqs::
Use this parameter to display a counter tracking requests that are sent to
```

```
the volume with fast-path.

``[-remote-reqs <integer>]`` - Number of Remote Reqs::
Use this parameter to display a counter tracking requests that are sent to
the volume with slow-path to remote Dblade using a CSM session.

``[-policy-name <text>]`` - Export Policy Name::
Use this parameter to display the Policy Name for this volume.

``[-trunking-status {true|false}]`` - Trunking Status::
Use this parameter to display trunking status for the specified
client/connection.

== Examples

The following example displays the connected clients information.
```

cluster-1::*> nfs connected-clients show
Node: node1
Vserver: vs1
Data-Ip: xxx.xxx.xxx.xxx
Local Remote
Client-Ip Protocol Volume Policy Idle-Time Reqs Reqs Trunking
--- --- --- --- --- --- --- ---
xxx.xxx.xxx.xxx nfs4.1 root_vol pol1 6s 9 0 false
xxx.xxx.xxx.xxx nfs3 vol1 pol1 56s 7 0 false
xxx.xxx.xxx.xxx nfs4.1 vol1 pol1 6s 7 0 false
3 entries were displayed.

```
[[ID01aa5fc7050565b898f9532df3d1d5ea]]
= vserver nfs credentials count
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Count credentials cached by NFS

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver nfs credentials count` command displays the number of
credentials stored in NFS credentials cache on a specific node. This
command has no effect if the specified node has no active data.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
The name of the node on which the command is executed.

== Examples

Lists the number of credentials stored by NFS on node node1
```

cluster1::*> vserver nfs credentials count -node node1

Number of credentials cached by NFS on node "node1": "2"

```
[[IDf99a147289e48f612fb9966b88f58674]]
= vserver nfs credentials flush
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Flush credentials cached by NFS

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.
```

```
== Description

The `vserver nfs credentials flush` command deletes credentials from the
NFS credentials cache on a specific node for a given Vserver or a given
UNIX user. This command has no effect if the vserver that is specified has
no active data interfaces on the node where the command is run.

== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to clear the credential cache for the Vserver you
specify.

{ ``[-unix-user-id <integer>]`` - UNIX User ID [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to clear the credential cache for the UNIX user id you
specify.

``[-client-ip <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to clear the credential cache for the client IP address
you specify.

| ``[-unix-user-name <text>]`` - UNIX User Name [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to clear the credential cache for the UNIX user name
you specify.

| ``[-unix-group-id <integer>]`` - UNIX Group ID [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to clear the credential cache for the UNIX group id you
specify.

| ``[-unix-group-name <text>]`` - UNIX Group Name [.advaced-
priv]#(privilege: advanced)# }::
Use this parameter to clear the credential cache for the UNIX group name
you specify.

== Examples

Clear the credential cache for user user1 on node node1 in Vserver vs1.
```

cluster1::*> vserver nfs credentials flush -node node1 -vserver vs1 -unix-user-name user1

Number of matching credentials flushed: 1

```
[[ID4aba186e1e0d464ab8b52e50b0c048b3]]
= vserver nfs credentials show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show credentials cached by NFS

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver nfs credentials show` command displays the user account
credentials stored on a specific node for a given UNIX user. This command
has no effect if the vserver specified has no active data interfaces on
the node where the command is run.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to search for user credentials in the Vserver you
specify.

{ ``-unix-user-id <integer>`` - UNIX User ID [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to search for user credentials for the UNIX user id you
specify.
```

```
``[-client-ip <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to search for user credentials for the client ip you
specify.

| ``-unix-user-name <text>`` - UNIX User Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to search for user credentials for the UNIX user name
you specify.

| ``-unix-group-id <integer>`` - UNIX Group ID [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to search for user credentials for the UNIX group id
you specify.

| ``-unix-group-name <text>`` - UNIX Group Name [.advaced-
priv]#(privilege: advanced)# }::
Use this parameter to search for user credentials for the UNIX group name
you specify.

``[-flags {ip-qualifier-configured|ip-qualifier-not-configured|unix-
extended-creds-present|no-unix-extended-creds|unix-extended-creds-
requested|unix-creds-transient-failure|cifs-creds-present|no-cifs-
creds|cifs-creds-requested|cifs-cifs-transient-failure|place-
holder|transient-failure|transient-error-on-last-refresh|id-name-mapping-
present|no-id-name-mapping|id-name-mapping-requested|id-name-mapping-
transient-failure|unix-cred-is-partial}]`` - Credential Entry Flags
[.advaced-priv]#(privilege: advanced)#::
The credential entry flags.

``[-last-refresh-time <[<integer>h][<integer>m][<integer>s]>]`` - Time
since Last Refresh [.advaced-priv]#(privilege: advanced)#::
Time since last refreshed.

``[-last-access-time <[<integer>h][<integer>m][<integer>s]>]`` - Time
since Last Access [.advaced-priv]#(privilege: advanced)#::
Time since last access.

``[-hit-count <integer>]`` - Number of Hits [.advaced-priv]#(privilege:
advanced)#::
Number of times the cached credential is fetched successfully.

``[-unix-cred-flags <integer>]`` - UNIX Credential Flags [.advaced-
priv]#(privilege: advanced)#::
UNIX credentials flags.
```

```
``[-unix-cred-domain-id <integer>]`` - UNIX Credential Domain ID
[.advaced-priv]#(privilege: advanced)#::
UNIX credentials domain ID.

``[-unix-cred-uid <integer>]`` - UNIX Credential UID [.advaced-
priv]#(privilege: advanced)#::
User ID of the UNIX user.

``[-unix-cred-primary-gid <integer>]`` - UNIX Credential Primay GID
[.advaced-priv]#(privilege: advanced)#::
Primary GID of the UNIX user.

``[-unix-cred-additional-gids <integer>,...]`` - UNIX Credential
Additional GIDs [.advaced-priv]#(privilege: advanced)#::
Additional GIDs of the UNIX user.

``[-win-cred-flags <integer>]`` - Windows Credential Flags [.advaced-
priv]#(privilege: advanced)#::
Windows credentials flags.

``[-win-cred-user-sid <text>]`` - Windows Credential User SID [.advaced-
priv]#(privilege: advanced)#::
SID of the windows user.

``[-win-cred-primary-group-sid <text>]`` - Windows Credential Primary
Group SID [.advaced-priv]#(privilege: advanced)#::
SID of the windows user's primary group.

``[-win-cred-domain-sids <text>,...]`` - Windows Credential Domain SIDs
[.advaced-priv]#(privilege: advanced)#::
Domain SIDs of the windows user.

== Examples

Show the credentials cached by NFS Cred Store for the UNIX user node1 on
node node1.
```

gnk1cluster1::*> vserver nfs credentials show -node gnk1cluster1-01 -vserver coke -unix-user-name root

Credentials

Node: gnk1cluster1-01
Vserver: coke
Client IP: -
Flags: ip-qualifier-not-configured, unix-extended-creds-present, cifs-creds-present, id-name-mapping-present
Time since Last Refresh: 10s
Time since Last Access: 5s
Hit Count: 24

Unix Credentials:
Flags: 0
Domain ID: 0
UID: 0
Primary GID: 1
Additional GIDs: 1
Windows Credentials:
Flags: 8759
User SID: S-1-5-21-2552784647-1202982559-4146209732-500
Primary Group SID: S-1-5-21-2552784647-1202982559-4146209732-513
Domain SIDs: S-1-5-21-2552784647-1202982559-4146209732
S-1-1
S-1-5
S-1-5-32
ID-Name Information:
Type: user
ID: 0
Name: root

```
[[IDe7723c3c02677bce294572520f7bbcbb]]
= vserver nfs kerberos interface disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable NFS Kerberos on a LIF

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos interface disable` command disables NFS Kerberos
on a logical interface.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver in which the logical interface
exists.

``-lif <lif-name>`` - Logical Interface::
This parameter specifies the name of the logical interface on which you
want to disable NFS Kerberos.

``[-admin-username <text>]`` - Account Creation Username::
This optional parameter specifies the administrator user name.

``[-admin-password <text>]`` - Account Creation Password::
This optional parameter specifies the administrator password.

== Examples

The following example disables NFS Kerberos on a Vserver named vs0 and a
logical interface named datalif1.
```

vs1::> vserver nfs kerberos interface disable -vserver vs0 -lif datalif1

```
[[ID64954592ba7fa5115279d679cb164a2c]]
= vserver nfs kerberos interface enable
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable NFS Kerberos on a LIF

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos interface enable` command enables NFS Kerberos
on a logical interface.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver in which the logical interface
exists.

``-lif <lif-name>`` - Logical Interface::
This parameter specifies the name of the logical interface on which you
want to enable NFS Kerberos.

``[-spn <text>]`` - Service Principal Name::
This optional parameter specifies the service principal name (SPN) for the
logical interface you want to enable. This value must be in the form
nfs/``_host_name_`` @``_REALM_`` , where ``_host_name_`` is the fully
qualified host name of the Kerberos server, nfs is the service, and
``_REALM_`` is the name of the Kerberos realm (for instance, EXAMPLE.COM).
Specify Kerberos realm name in uppercase.

``[-admin-username <text>]`` - Account Creation Username::
This optional parameter specifies the administrator user name.

``[-admin-password <text>]`` - Account Creation Password::
This optional parameter specifies the administrator password.

``[-keytab-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}]``
- Load Keytab from URI::
This optional parameter specifies loading a keytab file from the specified
URI.

``[-ou <text>]`` - Organizational Unit::
This optional parameter specifies the organizational unit (OU) under which
the Microsoft Active Directory server account will be created when you
enable Kerberos using a realm for Microsoft KDC. If this parameter is not
```

```
specified, the default OU is "CN=Computers".

``[-machine-account <text>]`` - Machine Account Name::
This optional parameter specifies the machine account to create in Active
Directory

== Examples

The following example enables NFS Kerberos on a Vserver named vs0 and a
logical interface named datalif1. The SPN is
nfs/sec.example.com@AUTH.SEC.EXAMPLE.COM and the keytab file is loaded
from ftp://ftp.example.com/keytab.
```

vs1::> vserver nfs kerberos interface enable -vserver vs0 -lif datalif1 -spn nfs/sec.example.com@AUTH.SEC.EXAMPLE.COM -keytab-uri ftp://ftp.example.com/keytab

```
[[IDbbb8784fc5aa7ed3f1f2a6525217fff3]]
= vserver nfs kerberos interface modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the Kerberos configuration of an NFS server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos interface modify` command modifies a Kerberos
configuration for NFS. An NFS Kerberos configuration is associated with
both a Vserver and a logical interface.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver associated with the NFS Kerberos
configuration you want to modify.

``-lif <text>`` - Logical Interface::
This parameter specifies the name of the logical interface associated with
the NFS Kerberos configuration you want to modify.
```

```
``[-kerberos {enabled|disabled}]`` - Kerberos Enabled::
This optional parameter specifies whether to enable or disable Kerberos
for NFS on the specified Vserver and logical interface. If you specify a
value of `enable` , you must also specify the -spn parameter. The command
prompts you for a user name and password for a Kerberos principal in the
same realm as the principal specified by the `-spn` parameter; this
principal must have permission to create or modify the principal specified
by the `-spn` parameter.

``[-spn <text>]`` - Service Principal Name::
This optional parameter specifies the service principal name (SPN) of the
Kerberos configuration you want to modify. If you specify a value of
`enable` for the `-kerberos` parameter, you must also specify this
parameter. This value must be in the form nfs/``_host_name_`` @``_REALM_``
, where ``_host_name_`` is the fully qualified host name of the Kerberos
server, nfs is the service, and ``_REALM_`` is the name of the Kerberos
realm (for instance, EXAMPLE.COM). Specify Kerberos realm names in
uppercase.

``[-admin-username <text>]`` - Account Creation Username::
This optional parameter specifies the administrator user name.

``[-keytab-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}]``
- Load Keytab from URI::
This optional parameter specifies loading a keytab file from the specified
URI.

``[-ou <text>]`` - Organizational Unit::
This optional parameter specifies the organizational unit (OU) under which
the Microsoft Active Directory server account will be created when you
enable Kerberos using a realm for Microsoft KDC. If this parameter is not
specified, the default OU is "CN=Computers".

``[-machine-account <text>]`` - Machine Account Name::
This optional parameter specifies the machine account to create in Active
Directory

== Examples

The following example enables an NFS Kerberos configuration on a Vserver
named vs0 and a logical interface named datalif1. The SPN is
nfs/sec.example.com@AUTH.SEC.EXAMPLE.COM and the keytab file is loaded
from ftp://ftp.example.com/keytab.
```

vs1::> vserver nfs kerberos interface modify -vserver vs0 -lif datalif1
-kerberos enabled -spn nfs/sec.example.com@AUTH.SEC.EXAMPLE.COM -keytab-uri

ftp://ftp.example.com/keytab

```
[[ID2a54b4848fade06ffb4cc06f4954de2c]]
= vserver nfs kerberos interface show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Kerberos configurations of NFS servers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos interface show` command displays information
about Kerberos configurations for NFS. The command output depends on the
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all
Kerberos configurations for NFS:

* Vserver name
* Logical interface name
* Logical interface IP address
* Whether Kerberos is enabled or disabled
* The Kerberos service principal name (SPN)
* The permitted encryption types

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
Kerberos configurations for NFS that are enabled, run the command with the
`-kerberos enabled` parameter.

== 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.
```

```
``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter and the `-lif` parameter, the command
displays information only about the Kerberos configuration or
configurations for NFS that are associated with the specified Vserver and
logical interface.

``[-lif <text>]`` - Logical Interface::
If you specify this parameter and the `-vserver` parameter, the command
displays information only about the Kerberos configuration or
configurations for NFS that are associated with the specified logical
interface and Vserver.

``[-address <IP Address>]`` - IP Address::
If you specify this parameter, the command displays information only about
the Kerberos configurations for NFS that are associated with the specified
logical-interface IP address.

``[-kerberos {enabled|disabled}]`` - Kerberos Enabled::
If you specify this parameter, the command displays information only about
the Kerberos configurations for NFS that match the specified value.

``[-spn <text>]`` - Service Principal Name::
If you specify this parameter, the command displays information only about
the Kerberos configuration or configurations for NFS that match the
specified SPN.

``[-permitted-enc-types <NFS Kerberos Encryption Type>,...]`` - Permitted
Encryption Types::
If you specify this parameter, the command displays information only about
the Kerberos configuration for NFS that matches the specified encryption
types.

``[-machine-account <text>]`` - Machine Account Name::
If you specify this parameter, the command displays information only about
the Kerberos configuration for NFS that matches the specified machine
account.

== Examples

The following example displays information about the Kerberos
configuration for NFS associated with the Vserver vs0 and the logical
interface datalif1:
```

vs1::> vserver nfs kerberos interface show -vserver vs0 -lif datalif1
Vserver: vs1
Logical Interface: datalif1

IP Address: 192.0.2.130
Kerberos Enabled: enabled
Service Principal Name: nfs/sec.example.com@AUTH.SEC.EXAMPLE.COM
Permitted Encryption Types: des,des3,aes-128,aes-256

```
[[IDbaad967e9dae944ec87eec1bae3301d7]]
= vserver nfs kerberos realm create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a Kerberos realm configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos realm create` command creates a Kerberos realm
configuration.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver associated with the Kerberos realm
configuration that you want to create.

``-realm <text>`` - Kerberos Realm::
This parameter specifies the name of the Kerberos realm for the
configuration.

``-kdc-vendor <Kerberos Key Distribution Center (KDC) Vendor>`` - KDC
Vendor::
This optional parameter specifies the KDC vendor. Specify Microsoft if you
are using a Microsoft Active Directory server; specify Other if you are
using a UNIX server.

``-kdc-ip <IP Address>`` - KDC IP Address::
This parameter specifies the IP address of the Kerberos Distribution
Center (KDC) server.

``[-kdc-port <integer>]`` - KDC Port::
This optional parameter specifies the port number of the KDC server. The
default setting is 88.
```

```
``[-clock-skew <integer>]`` - Clock Skew::
This optional parameter specifies how many minutes of clock skew between
the clients and the server are permitted. The default setting is 5
minutes.

``[-adserver-name <text>]`` - Active Directory Server Name::
This parameter specifies the name of an Active Directory server for the
configuration. Use this parameter only if you specified the value of `-
kdc-vendor` parameter as ``_Microsoft_`` .

``[-adserver-ip <IP Address>]`` - Active Directory Server IP Address::
This parameter specifies the IP address of an Active Directory server for
the configuration. Use this parameter only if you specified the value of
the `-kdc-vendor` parameter as ``_Microsoft_`` .

``[-comment <text>]`` - Comment::
This optional parameter specifies a comment for the Kerberos realm
configuration.

``[-adminserver-ip <IP Address>]`` - Admin Server IP Address::
This optional parameter specifies the IP address of the administrative
server. Use this parameter only if you specified the value of `-kdc-
vendor` parameter as Other. The default setting for this parameter is the
KDC server's IP address as specified by the `-kdc-ip` parameter.

``[-adminserver-port <integer>]`` - Admin Server Port::
This optional parameter specifies the port number of the administrative
server. The default setting is 749. Use this parameter only if you
specified the value of `-kdc-vendor` parameter as Other.

``[-passwordserver-ip <IP Address>]`` - Password Server IP Address::
This optional parameter specifies the IP address of the password server.
Use this parameter only if you specified the value of `-kdc-vendor`
parameter as Other. The default setting for this parameter is the KDC
server's IP address as specified by the `-kdc-ip` parameter.

``[-passwordserver-port <integer>]`` - Password Server Port::
This optional parameter specifies the port number of the password server.
The default setting is 464. Use this parameter only if you specified the
value of `-kdc-vendor` parameter as Other.

== Examples

The following example creates a Kerberos realm named SEC.EXAMPLE.COM for
the Vserver named AUTH. The permitted clock skew is 15 seconds. The KDC's
IP address is 192.0.2.170 and its port is 88. The KDC vendor is Other (for
```

```
a UNIX KDC). The administrative server's IP address is 192.0.2.170 and its
port is 749. The password server's IP address is 192.0.2.170 and its port
is 464.
```

cluster1::> vserver nfs kerberos realm create -vserver AUTH -realm SEC.EXAMPLE.COM -clock-skew 15 -kdc -ip 192.0.2.170 -kdc-port 88 -kdc-vendor Other -adminserver-ip 192.0.2.170 -adminserver-port 749 -passwordserver-ip 192.0.2.170 -passwordserver-port 464

```
[[IDb028e8e002a499b170719fa87ec61ea4]]
= vserver nfs kerberos realm delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Kerberos realm configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos realm delete` command deletes a Kerberos realm
configuration from the system.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for the Kerberos realm
configuration that you want to delete.

``-realm <text>`` - Kerberos Realm::
This parameter specifies the name of the Kerberos realm for the
configuration.

== Examples

The following example deletes the Kerberos realm SEC.EXAMPLE.COM from the
Vserver named AUTH:
```

cluster1::> vserver nfs kerberos realm delete -vserver AUTH -realm SEC.EXAMPLE.COM

```
[[IDecbc90b2a9811efb8cfe93170c3f5c35]]
= vserver nfs kerberos realm modify
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a Kerberos realm configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos realm modify` command modifies one or more
attributes of a Kerberos realm configuration.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for the Kerberos realm
configuration that you want to modify.

``-realm <text>`` - Kerberos Realm::
This optional parameter specifies the name of a Kerberos realm for the
configuration.

``[-kdc-vendor <Kerberos Key Distribution Center (KDC) Vendor>]`` - KDC
Vendor::
This optional parameter specifies the KDC vendor. Specify Microsoft if you
are using a Microsoft Active Directory server; specify Other if you are
using a UNIX server.

``[-kdc-ip <IP Address>]`` - KDC IP Address::
This parameter specifies the IP address of the Kerberos Distribution
Center (KDC) server.

``[-kdc-port <integer>]`` - KDC Port::
This optional parameter specifies the port number of the KDC server. The
default setting at the time of creation is 88.

``[-clock-skew <integer>]`` - Clock Skew::
This optional parameter specifies how many minutes of clock-skew between
server and the clients are permitted. The default setting at the time of
creation is 5 minutes.

``[-adserver-name <text>]`` - Active Directory Server Name::
This parameter specifies the name of an Active Directory server for the
```

```
configuration. Use this parameter if you specified the value of `-kdc-
vendor` parameter as Microsoft.

``[-adserver-ip <IP Address>]`` - Active Directory Server IP Address::
This parameter specifies the IP address of an Active Directory server for
the configuration. Use this parameter if you specified the value of the `-
kdc-vendor` parameter as Microsoft.

``[-comment <text>]`` - Comment::
This optional parameter specifies a comment for the Kerberos realm
configuration.

``[-adminserver-ip <IP Address>]`` - Admin Server IP Address::
This optional parameter specifies the IP address of the administrative
server. Use this parameter if you specified the value of `-kdc-vendor`
parameter as Other.

``[-adminserver-port <integer>]`` - Admin Server Port::
This optional parameter specifies the port number of the administrative
server. The default setting at the time of creation is 749. Use this
parameter if you specified the value of the `-kdc-vendor` parameter as
Other.

``[-passwordserver-ip <IP Address>]`` - Password Server IP Address::
This optional parameter specifies the IP address of the password server.
Use this parameter if you specified the value of `-kdc-vendor` parameter
as Other.

``[-passwordserver-port <integer>]`` - Password Server Port::
This optional parameter specifies the port number of the password server.
The default setting at the time of creation is 464. Use this parameter
only if you specified the value of `-kdc-vendor` parameter as Other.

== Examples

The following example modifies the Kerberos realm SEC.EXAMPLE.COM for the
Vserver named AUTH to use a Microsoft KDC server with the IP address
192.0.2.170 and an Active Directory server named AUTH.SEC.EXAMPLE.COM with
the IP address 192.0.2.170:
```

cluster1::> vserver nfs kerberos realm modify -vserver AUTH -realm SEC.EXAMPLE.COM -adserver-name AUTH.SEC.EXAMPLE.COM -adserver-ip 192.0.2.170 -kdc-ip 192.0.2.170 -kdc-vendor Microsoft

```
[[IDf7d7520ffeb7e08a8c38377aab25679a]]
= vserver nfs kerberos realm show
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Kerberos realm configurations

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs kerberos realm show` command displays information about
Kerberos realm configurations. The command output depends on the
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all
Kerberos realm configurations:

* Vserver
* Kerberos realm name
* Active Directory server name
* Kerberos Distribution Center (KDC) vendor
* KDC IP address
* The permitted encryption types



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations for the specified Vserver.

``[-realm <text>]`` - Kerberos Realm::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified Kerberos realm.

``[-kdc-vendor <Kerberos Key Distribution Center (KDC) Vendor>]`` - KDC
```

```
Vendor::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified KDC vendor.

``[-kdc-ip <IP Address>]`` - KDC IP Address::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified KDC IP address.

``[-kdc-port <integer>]`` - KDC Port::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified KDC port number.

``[-clock-skew <integer>]`` - Clock Skew::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified clock skew.

``[-adserver-name <text>]`` - Active Directory Server Name::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the Active Directory server
that has the specified name.

``[-adserver-ip <IP Address>]`` - Active Directory Server IP Address::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the Active Directory server
that has the specified IP address.

``[-comment <text>]`` - Comment::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that match the specified comment text.

``[-adminserver-ip <IP Address>]`` - Admin Server IP Address::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified administrative-
server IP address.

``[-adminserver-port <integer>]`` - Admin Server Port::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified administrative-
server port number.

``[-passwordserver-ip <IP Address>]`` - Password Server IP Address::
If you specify this parameter, the command displays information only about
the Kerberos realm configurations that use the specified password-server
IP address.

``[-passwordserver-port <integer>]`` - Password Server Port::
If you specify this parameter, the command displays information only about
```

```
the Kerberos realm configurations that use the specified password-server
port number.

``[-permitted-enc-types <NFS Kerberos Encryption Type>,...]`` - Permitted
Encryption Types::
If you specify this parameter, the command displays information only about
the Kerberos realm configuration that match the specified encryption
types.

== Examples

The following example displays information about all Kerberos realm
configurations:
```

cluster1::> vserver nfs kerberos realm show
Kerberos Active Directory KDC KDC
Vserver Realm Server Vendor IP Address
-------- ---------- ---------- ---------- ----------
AUTH SEC.EXAMPLE.COM AUTH.SEC.EXAMPLE.COM
Microsoft 192.0.2.170

```
[[ID0e4955316ab43d9d1b2c185d2aad3a7c]]
= vserver nfs pnfs devices create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new pNFS device and its mapping

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

 The `vserver nfs pnfs devices create` command creates a pNFS device for a
given instance of a volume. The actual creation of pNFS devices is
automatically done by the pNFS implementation in ONTAP kernel. The usage
of this command might interfere with the functionality of the pNFS server.
Thus, it is advised that this command not be used without supervision by
technical support.

== Parameters

``[-global-device-table-id <integer>]`` - Global Device Mapping Table ID
```

```
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies the unique identifier that the pNFS
devices subsystem assigns to the device that corresponds to the MSID
described below. The pNFS devices implementations keeps track of the
global unique identifier that needs to be assigned to this device. It is
expected that users need not specifically input the device identifier
while creation.

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver to which the volumes belong.

``-msid <integer>`` - Volume MSID [.advaced-priv]#(privilege: advanced)#::
This parameter uniquely identifies the volume for which you are creating a
pNFS device.

``-striping-epoch <integer>`` - Striping Epoch [.advaced-priv]#(privilege:
advanced)#::
This optional parameter specifies the striping epoch identifier for a
volume for which you are creating a pNFS device. For flexible volumes, the
value is always 1.

``-device-access <integer>`` - Device Access Flags [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the type of access that is given to the
pNFS device that you are creating. If the value is 1, it means write
access. If the value if 0, it means read access. By default, the device is
created with write access.

``-version <integer>`` - Device Version [.advaced-priv]#(privilege:
advanced)#::
This optional parameter specifies the version associated with the pNFS
device identifier. By default, the version is set to 1.

``[-generation-count <integer>]`` - Device Generation [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the generation count associated with the
pNFS device identifier. If a device already exists, the existing device is
invalidated and the generation number for the device is bumped. If a
device does not already exist, a new device is created with generation
number 1.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Device Creation Time [.advaced-
priv]#(privilege: advanced)#::
This optional parameter specifies the time at which the device is created.
If the parameter is not specified, the time at which the device is created
```

```
is stored along with the device.

``[-mapping-status {available|notavailable}]`` - Device Mapping Status
[.advaced-priv]#(privilege: advanced)#::
This optional parameter specifies if the mapping exists for a device. If
the value is set to "available", the mappings will be created in the
device mappings table. If the value is set to "notavailable", the mappings
will not be created in the device mappings table.

== Examples





[[ID4213640048bcd6910e1c73be4a730812]]
= vserver nfs pnfs devices delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a pNFS device

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver nfs pnfs devices delete` command deletes a unique pNFS
device. The pNFS device to be deleted is identified by the unique device
mapping identifier (mid) parameter passed to this operation. When this
operation is successful, the device mappings corresponding to the device
and the information corresponding to the device itself are removed. You
can obtain the global mapping identifier from the list of devices using
the command xref:{relative_path}vserver-nfs-pnfs-devices-show.html[vserver
nfs pnfs devices show] .

== Parameters

``-global-device-table-id <integer>`` - Global Device Mapping Table ID
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the pNFS global device mapping identifier that
uniquely identifies a pNFS device
```

```
== Examples

The following example deletes the device information of a device with
global mapping identifier value 2.
```

cluster1::> vserver nfs pnfs delete -mid 2

```
== Related Links

* xref:{relative_path}vserver-nfs-pnfs-devices-show.html[vserver nfs pnfs
devices show]




[[ID5baa014bd160332e177789523c896379]]
= vserver nfs pnfs devices show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display pNFS device information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

 The `vserver nfs pnfs devices show` command displays a pNFS device for a
given instance of a volume. The command output depends on the parameter or
parameters specified with the command. If you do not specify any
parameters, the command displays the following information about all pNFS
devices:

* Vserver name
* The global device mapping identifier of the device
* The master data set ID (MSID) of the volume that leads to this device
* The mapping status of the device
* The generation number of the device

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
devices that are exported as write-only devices, enter the command with
the `-access-flags 1` parameter.
```

```
== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-global-device-table-id <integer>]`` - Global Device Mapping Table ID
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the unique identifier that the pNFS devices subsystem assigns to the
device that is being output.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the Vserver that owns the volume represented by MSID.

``[-msid <integer>]`` - Volume MSID [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the volume or volumes that match the specified MSID.

``[-striping-epoch <integer>]`` - Striping Epoch [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the striping epoch identifier for a volume that serves as the basis for
the pNFS device.

``[-device-access <integer>]`` - Device Access Flags [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
access flags which specify the type of access that is given to the pNFS
device. If the value is 1, it means write access. If the value is 0, it
means read access.

``[-version <integer>]`` - Device Version [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
pNFS devices that match the specified version number.

``[-generation-count <integer>]`` - Device Generation [.advaced-
```

```
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
generation count associated with the pNFS device identifier.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Device Creation Time [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
pNFS devices that were created at the specified time.

``[-mapping-status {available|notavailable}]`` - Device Mapping Status
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
if the mapping exists for a device. If the value is set to "available",
the mappings can be seen in the device mappings table. If the value is set
to "notavailable", the mappings will not be seen in the device mappings
table.

== Examples

The following example displays the information of a device with global
mapping identifier 6. The device corresponds to a volume with MSID
2147484673 on Vserver vs1. The device mappings corresponding to this
device follow in the mappings table.
```

cluster1::*> vserver nfs pnfs devices show
Vserver Name Mapping ID Msid Mapping Status Generation
--- --- --- --- ---
vs1 1 2147484673 available 6

cluster1::*> vserver nfs pnfs devices mappings show
Vserver Name Mapping ID Dsid LIF IP
--- --- --- ---
vs1 1 1025 10.53.4.14

```
[[ID9d4750af3f551a3108193afffd69d144]]
= vserver nfs pnfs devices cache show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the device cache

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.
```

```
== Description

 The `vserver nfs pnfs devices cache show` command displays the device
cache.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the pNFS devices cache present on the node.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the Vserver that has the pNFS devices cache.

== Examples





[[ID21e23c884d373fcebd74c011a1469e83]]
= vserver nfs pnfs devices mappings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the list of pNFS device mappings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description
```

```
 The xref:{relative_path}vserver-nfs-pnfs-devices-show.html[vserver nfs
pnfs devices show] command displays a pNFS device for a given instance of
a volume. The command output depends on the parameter or parameters
specified with the command. If you do not specify parameters, the command
displays the following information about all pNFS devices:

* Vserver name
* The global device mapping identifier of the device
* The Data Set ID (DSID) of the constituent volume
* The LIF IP address that serves the constituent on the same controller.

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
devices that are exported as write-only devices, enter the command with
the `-access-flags 1` parameter.

== Parameters

``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-global-device-table-id <integer>]`` - Global Device Mapping Table ID
[.advaced-priv]#(privilege: advanced)#::
This specifies the unique identifier that the pNFS devices subsystem
assigns to the device whose mappings are being output.

``[-vserver <vserver name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If your specify this parameter, the command displays information only
about the Vserver that the mapping identifier and DSID belong to.

``[-dsid <integer>]`` - Constituent Volume DSID [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the volume or volumes that match the specified DSID.

``[-lifip <IP Address>]`` - LIF IP Address [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the pNFS devices that match the specified LIF IP address.

== Examples
```

```
The following example displays the device information of a device with
global mapping identifier 6. The device corresponds to a volume with MSID
2147484673 on Vserver vs1. The device has one constituent with DSID 1025
and is served by the LIF with the IP address 10.53.4.14.
```

cluster1::*> vserver nfs pnfs devices* show
Vserver Name Mapping ID Msid Mapping Status Generation
--- --- --- --- ---
vs1 1 2147484673 available 6

cluster1::*> vserver nfs pnfs devices mappings show
Vserver Name Mapping ID Dsid Lif IP
--- --- --- ---
vs1 1 1025 10.53.4.14

```
== Related Links

* xref:{relative_path}vserver-nfs-pnfs-devices-show.html[vserver nfs pnfs
devices show]




[[IDaefc6e34f3739c44a7cdfb1c44903eea]]
= vserver nfs storepool show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display storepool information for currently connected NFSv4 clients

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver nfs storepool show` command shows the number of storepool
resources consumed. A storepool is a pool of resources used by NFSv4
clients. The command output depends on the parameter or parameters
specified with the command. If no parameters are specified, the command
displays the following storepool information about all the connected
clients:

* Node on which the storepool resource is located
```

```
* Vserver name
* Data ip address
* Client ip address
* Protocol (nfs4, nfs4.1, or nfs4.2)
* Session trunking status (true or false)
* Owner count
* Open state count
* Delegation state count
* Lock state count
* Bad State Count
* Stale State Count



To display detailed information about a single instance, run the command
with the `-instance` parameter. The detailed view provides all of the
information in the previous list and the following additional information:

* Layout state count
* Copy state count
* Client identifier



== Parameters

``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed. The fields Vserver,
node, data-lif-ip, client-ip and protocol are the default fields (see
example).

| ``[-instance  ]`` [.advaced-priv]#(privilege: advanced)# }::
If this parameter is specified, the command displays information about all
entries.

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified node.

``[-vserver <vserver>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified Vserver.

``[-data-lif-ip <IP Address>]`` - Data LIF IP Address [.advaced-
priv]#(privilege: advanced)#::
```

```
If this parameter is specified, the command displays information only
about clients with the specified data LIF.

``[-client-ip <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified client IP.

``[-protocol <Nfs4 Storepool Client Access Protocol>]`` - Protocol Version
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified protocol.

``[-trunking-status {true|false}]`` - Trunking Status [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified trunking status.

``[-owner-count <integer>]`` - Allocated Owner Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of owner counts.

``[-open-state-count <integer>]`` - Allocated OpenState Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of open state counts.

``[-deleg-state-count <integer>]`` - Allocated DelegState Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of delegation state counts.

``[-lock-state-count <integer>]`` - Allocated LockState Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of lock state counts.

``[-layout-state-count <integer>]`` - Allocated LayoutState Count
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of layout state counts.

``[-copy-state-count <integer>]`` - Allocated CopyState Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of copy state counts.
```

```
``[-client-id <integer>]`` - Unique Client Identifier [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified client id.

``[-exhausted-object {owner|open|lock|delegation}]`` - Exhausted Object
List [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified exhausted object.

``[-is-exhausted {true|false}]`` - Is Storepool Exhausted? [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified exhausted status.

``[-bad-seqid-state-count <integer>]`` - Bad Sequential ID State Count
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of bad sequential ID state counts.

``[-stale-stateid-count <integer>]`` - Stale State ID Count [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about clients with a matching number of stale state ID counts.

== Examples

The following example displays the storepool information of NFSv4 clients.
```

```
cluster1::*> vserver nfs storepool show
Node: node1
Vserver: vs1
Data-Ip: 10.0.1.1
Client-Ip Protocol IsTrunked OwnerCount OpenCount DelegCount LockCount
-------------- --------- --------- ---------- ---------- ---------- ---------
10.0.2.1 nfs4.1 true 2 1 0 4
10.0.2.2 nfs4.2 true 2 1 0 4
2 entries were displayed.

cluster1::*> vserver nfs storepool show -instance
Node Name: node1
Vserver: vs1
Data LIF IP Address: 10.0.1.1
Client IP Address: 10.0.2.1
Protocol Version: nfs4.1
Trunking Status: true
Allocated Owner Count: 2
Allocated OpenState Count: 1
```

Allocated DelegState Count: 0
Allocated LockState Count: 4
Allocated LayoutState Count: 0
Allocated CopyState Count: 0
Unique Client Identifier: 16962161483653840897
Exhausted Object List: -
Is Storepool Exhausted? : false
Bad Sequential ID State Count: 0
Stale State ID Count: 0
Node Name: node1
Vserver: vs1
Data LIF IP Address: 10.0.1.1
Client IP Address: 10.0.2.2
Protocol Version: nfs4.2
Trunking Status: true
Allocated Owner Count: 2
Allocated OpenState Count: 1
Allocated DelegState Count: 0
Allocated LockState Count: 4
Allocated LayoutState Count: 0
Allocated CopyState Count: 0
Unique Client Identifier: 16962161483653840897
Exhausted Object List: -
Is Storepool Exhausted? : false
Bad Sequential ID State Count: 0
Stale State ID Count: 0

2 entries were displayed.

```
[[ID9a7d8c0758e9cd36a7e80154be1c2c17]]
= vserver nfs storepool blocked-client flush
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Flush cached blocked client entries on the specified node

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver nfs storepool blocked-client flush` command flushes the
storepool blocked client cache which allows the previously blocked
specified clients to reconnect.

* Node on which storepool blocked client cache is present
* Client ip address



== Parameters

``-node <nodename>`` - Node [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command is used to flush the storepool
blocked client cache only in the specified node.

``[-client-ip <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command is used to only flush the
storepool blocked client cache associated with specified client IP.

== Examples

The following example flushes all NFSv4 storepool blocked client caches
present on node1.
```

cluster1::*> vserver nfs storepool blocked-client flush -node node1
Number of NFSv4.x blocked clients flushed on node "node1": 0.
cluster1::*>

```
[[ID353d1082a54635afbe7f7ed7c638f01b]]
```

```
= vserver nfs storepool blocked-client show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the Storepool Blocked Clients Information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver nfs storepool blocked-client show` command shows the NFSv4
storepool blocked clients. The command output depends on the parameter or
parameters specified with the command. If no parameters are specified, the
command displays the following information about all the blocked clients:

* Node on which storepool blocked client cache is present
* Client ip address
* Time to live for blocked client cache entry



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about clients with the specified node.

``[-client-ip <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information with the
specified client IP.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - TTL For
```

```
Entries (Secs) [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays all the blocked
client information whose time to live matches the provided value.

== Examples

The following example displays all NFSv4 storepool blocked clients caches
which are present on node1.
```

cluster1::*> vserver nfs storepool blocked-client show -node node1
cluster1::*>
Node: node1
Client-IP TTL (secs)
--- ---
1.0.1.62 512s
1.0.1.63 512s
1.0.1.126 512s
3 entries were displayed.
cluster1::*>

```
[[ID103f00b28c2b0be383d60a5f7dc5fb3f]]
= vserver nfs storepool config modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the storepool configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver nfs storepool config modify` command modifies NFSv4.x
storepool configuration parameters:

== Parameters

``[-limit-enforce {enabled|disabled}]`` - Storepool Limit Enforce Control
[.advaced-priv]#(privilege: advanced)#::
This parameter is used to enable or disable the limit enforcement. The
default value is ``_disabled_`` .

``[-blocked-client-ttl <integer>]`` - Blocked Client TTL in Seconds
[.advaced-priv]#(privilege: advanced)#::
This parameter is used to set the blocked client eviction time-to-live
(TTL). The client will be able to reconnect with NFs4.x server after this
value expires. Valid values are from 0 to 604800 seconds (7 days). The
default value is ``_86400_`` .

== Examples

The following is an example of how to modify the storepool limit-enforce
parameter.
```

cluster1::*> vserver nfs storepool config modify -limit-enforce disabled -blocked-client-ttl 90000

```
[[ID5afb80405fdfbada09acc7f9f0ed34d9]]
= vserver nfs storepool config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the storepool configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver nfs storepool config show` command displays the following
NFSv4.X storepool configuration information:

* Storepool limit enforce control
* Blocked client TTL in seconds



== Examples

The following example displays the NFSv4.x storepool configuration
information.
```

cluster1::*> vserver nfs storepool config show

Storepool Limit Enforce Control: disabled
Blocked client TTL in seconds: 90000

cluster1::*>

```
[[ID58ec53fc995d3fec8854266872120215]]
= vserver nfs tls interface disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable NFS TLS on a LIF

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs tls interface disable` command disables NFS TLS on a
logical interface.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver in which the logical interface
exists.

``-lif <lif-name>`` - Logical Interface::
This parameter specifies the name of the logical interface on which you
want to disable NFS TLS.

== Examples

The following example disables NFS TLS on a Vserver named vs0 and a
logical interface named datalif1.
```

vs1::> vserver nfs tls interface disable -vserver vs0 -lif datalif1

```
[[ID898bd6ad99b3c32062d3856d74c50f12]]
= vserver nfs tls interface enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable NFS TLS on a LIF

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs tls interface enable` command enables NFS TLS on a
logical interface.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver in which the logical interface
exists.

``-lif <lif-name>`` - Logical Interface::
This parameter specifies the name of the logical interface on which you
want to enable NFS TLS.

``[-certificate-name <text>]`` - TLS Certificate Name::
This parameter specifies the name of a certificate to be associated with
the instance of a given Vserver and logical interface.
+
NOTE: The use of self-signed SSL certificates exposes users to man-in-the-
middle security attacks. Where possible, obtain a certificate that is
signed by a reputable certificate authority (CA) and use the
xref:{relative_path}security-certificate-install.html[security certificate
install] command to configure it before enabling TLS on a Vserver and LIF.
+


== Examples

The following example enables NFS TLS on a Vserver named vs0 and a logical
interface named datalif1. The certificate name is lif1.com which is
preinstalled on server.
```

vs1::> vserver nfs tls interface enable -vserver vs0 -lif datalif1 -certificate-name lif1.com

```
== Related Links

* xref:{relative_path}security-certificate-install.html[security
certificate install]




[[ID8a905ea2d27f45d60995c08790d3c935]]
= vserver nfs tls interface modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the TLS configuration of an NFS server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs tls interface modify` command modifies a TLS
configuration for NFS. An NFS TLS configuration is associated with both a
Vserver and a logical interface.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver associated with the NFS TLS
configuration you want to modify.

``-lif <text>`` - Logical Interface::
This parameter specifies the name of the logical interface associated with
the NFS TLS configuration you want to modify.

``[-status {enabled|disabled}]`` - TLS Status::
This optional parameter specifies whether to enable or disable TLS for NFS
on the specified Vserver and logical interface. If you specify a value of
`enable` , you must also specify the -certificate-name parameter.

``[-certificate-name <text>]`` - TLS Certificate Name::
This optional parameter specifies the name of a certificate to be
```

```
associated with the instance of a given Vserver and logical interface. If
you specify a value of `enable` for the `-status` parameter, you must also
specify this parameter.
+
NOTE: The use of self-signed SSL certificates exposes users to man-in-the-
middle security attacks. Where possible, obtain a certificate that is
signed by a reputable certificate authority (CA) and use the
xref:{relative_path}security-certificate-install.html[security certificate
install] command to configure it before enabling TLS on a Vserver and LIF.
+


== Examples

The following example enables the NFS TLS configuration on a Vserver named
vs0 and a logical interface named datalif1. The certificate-name is
datalif1.example.com
```

clus1::> vserver nfs tls interface modify -vserver vs0 -lif datalif1 -status enabled -certificate-name datalif1.example.com

```
== Related Links

* xref:{relative_path}security-certificate-install.html[security
certificate install]




[[IDa1ebc121b6c53430e3f4d568d4881c4c]]
= vserver nfs tls interface show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the TLS configurations of NFS servers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nfs tls interface show` command displays information about
TLS configurations for NFS. The command output depends on the parameters
```

```
specified with the command. If you do not specify any parameters, the
command displays the following information about all TLS configurations
for NFS:

* Vserver name
* Logical interface name
* Logical interface IP address
* Whether TLS is enabled or disabled
* The TLS certificate name
* The TLS certificate UUID

You can specify additional parameters to display only information that
matches those parameters. For instance, to display information only about
TLS configurations for NFS that are enabled, run the command with the `-
status enabled` parameter.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter and the `-lif` parameter, the command
displays information only about the TLS configuration or configurations
for NFS that are associated with the specified Vserver and logical
interface.

``[-lif <text>]`` - Logical Interface::
If you specify this parameter and the `-vserver` parameter, the command
displays information only about the TLS configuration or configurations
for NFS that are associated with the specified logical interface and
Vserver.

``[-address <IP Address>]`` - IP Address::
If you specify this parameter, the command displays information only about
the TLS configurations for NFS that are associated with the specified
logical-interface IP address.

``[-status {enabled|disabled}]`` - TLS Status::
If you specify this parameter, the command displays information only about
the TLS configurations for NFS that match the specified value.
```

```
``[-certificate-name <text>]`` - TLS Certificate Name::
If you specify this parameter, the command displays information only about
the TLS configurations for NFS that match the specified value.

``[-certificate-uuid <UUID>]`` - TLS Certificate UUID::
If you specify this parameter, the command displays information only about
the TLS configuration for NFS that match the specified value.

== Examples

The following example displays information about the TLS configuration for
NFS associated with the Vserver vs0 and the logical interface datalif1:
```

vs1::> vserver nfs tls interface show -vserver vs0 -lif datalif1
Vserver: vs0
Logical Interface: datalif1
IP Address: 192.0.2.130
TLS Status: enabled
TLS Certificate Name: vs0_17A014E1FA994582
TLS Certificate UUID: d1905777-98e8-11ee-aa96-005056bb844f

```
:leveloffset: -1


= vserver nvme commands

:leveloffset: +1


[[ID65404bc18d61049267e35d50ab689b49]]
= vserver nvme create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create NVMe service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme create` command creates an NVMe service for a Vserver.

When you create the NVMe service on a Vserver, the Vserver must have only
``_nvme_`` in the allowed-protocols list.

When you create the NVMe service on a Vserver, the administrative status
of the service is ``_up_`` by default.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the NVMe service.

``[-status-admin {down|up}]`` - Administrative Status::
Specifies the configured administrative status of a service.

``[-mdns-service-discovery-enabled {true|false}]`` - Use Multicast DNS
Service Discovery [.advaced-priv]#(privilege: advanced)#::
Specifies the configured mDNS service discovery status of a service. If
you set this parameter to ``_false_`` , the Vserver will not advertise
NVMe-IP LIFs over Multicast DNS. The default value is ``_true_`` .

== Examples
```

cluster1::*> vserver nvme create -vserver vs_1

```
Creates an NVMe service on Vserver ``_vs_1_`` .




[[IDc002e88291c7816f406029b8b7b506be]]
= vserver nvme delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete NVMe service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme delete` command deletes an NVMe service of a Vserver.
Before you can delete an NVMe service, the administrative status must be
``_down_`` . Use the xref:{relative_path}vserver-nvme-modify.html[vserver
nvme modify] command to change the administrative status.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the NVMe service.

== Examples
```

cluster1::*> vserver nvme delete -vserver vs_1

```
Deletes the NVMe service on Vserver ``_vs_1_`` .

== Related Links

* xref:{relative_path}vserver-nvme-modify.html[vserver nvme modify]




[[ID46b8cb74eef4e4fa3562e3429fd3a9a1]]
= vserver nvme modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify NVMe service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme modify` command modifies an NVMe service configuration
on a Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver for the NVMe service.

``[-status-admin {down|up}]`` - Administrative Status::
Specifies the configured administrative status of a service. If you set
this parameter to ``_down_`` , the Vserver will not serve NVMe traffic.

``[-mdns-service-discovery-enabled {true|false}]`` - Use Multicast DNS
Service Discovery [.advaced-priv]#(privilege: advanced)#::
Specifies the configured mDNS service discovery status of a service. If
you set this parameter to ``_false_`` , the Vserver will not advertise
NVMe-IP LIFs over Multicast DNS. The default value for create is
``_true_`` .

== Examples
```

cluster1::*> vserver nvme modify -vserver vs_1 -status-admin down

```
Sets the administrative status of the NVMe service on Vserver ``_vs_1_``
to ``_down_`` .




[[ID83a6601abce8ae3953d8bc0cedc2c421]]
= vserver nvme show-discovery-controller
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display active NVMe discovery controllers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme show-discovery-controller` command displays information
for established NVMe controllers on the discovery subsystem. An NVMe
controller is established upon each host connection to the discovery
subsystem.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the NVMe controllers that match the Vserver
that you specify.

``[-controller-id <Hex 16bit Integer>]`` - Controller ID::
Use this parameter to display the NVMe controllers that match the
controller ID that you specify. The controller ID is the identifier
```

```
assigned by the NVMe subsystem upon host login and controller creation.

``[-discovery-nqn <text>]`` - Discovery Subsystem NQN::
Use this parameter to display the NVMe controllers that match the
discovery subsystem NQN that you specify.

``[-lif <text>]`` - Logical Interface::
Use this parameter to display the NVMe controllers that match the LIF that
you specify.

``[-node <nodename>]`` - Node::
Use this parameter to display the NVMe controllers that match the cluster
node that you specify.

``[-host-nqn <text>]`` - Host NQN::
Use this parameter to display the NVMe controllers that match the host NQN
that you specify. The host NQN is the NVMe identifier assigned to a
specific host.

``[-transport-protocol {fc-nvme|nvme-tcp}]`` - Transport Protocol::
Use this parameter to display the NVMe controllers that match the
transport protocol that you specify.

``[-initiator-transport-address <text>]`` - Initiator Transport Address::
Use this parameter to display the NVMe controllers that match the
initiator transport address that you specify. The initiator transport
address format depends on the transport protocol in use.

``[-trsvcid <text>]`` - Transport Service Identifier::
Use this parameter to display the NVMe controllers that match the
transport service identifier that you specify.

``[-host-id <Hex String>]`` - Host Identifier::
Use this parameter to display the NVMe controllers that match the host
identifier that you specify. The host identifier is a 128-bit identifier
assigned to a specific host.

``[-admin-queue-depth <integer>]`` - Admin Queue Depth::
Use this parameter to display the NVMe controllers that match the
administrative queue depth that you specify.

``[-header-digest-enabled {true|false}]`` - Header Digest Enabled::
Use this parameter to display the NVMe controllers that have header digest
enabled.

``[-data-digest-enabled {true|false}]`` - Data Digest Enabled::
Use this parameter to display the NVMe controllers that have data digest
```

```
enabled.

``[-kato <integer>]`` - Keep-Alive Timeout (msec)::
Use this parameter to display the NVMe controllers that match the keep-
alive timeout value (in msec) that you specify.

== Examples
```

cluster1::*> vserver nvme show-discovery-controller -vserver vs1
Vserver ID LIF Host NQN
------- ----- ---- ----------
vs1 0041h node1_e0f
nqn.2001-08.example.com:nvme:host1
vs1 0080h node2_e0f
nqn.2001-08.example.com:nvme:host2
2 entries were displayed.

```
[[ID54fba0d7007630cdb2d576e6ed5d5f63]]
= vserver nvme show-host-priority
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NVMe Host Priority Information

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver nvme show-host-priority` command displays NVMe I/O queue
count and NVMe I/O queue depth for each host priority level and node in
the cluster.

== 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>|local}]`` - Node::
Use this parameter to display the NVMe priority, NVMe I/O queue count, and
NVMe I/O queue depth that match the cluster node that you specify.

``[-transport-protocol {fc-nvme|nvme-tcp}]`` - Transport Protocol::
Use this parameter to display the priority, NVMe I/O queue count, and NVMe
I/O queue depth that match the transport protocol that you specify.

``[-priority {regular|high}]`` - Host Priority::
Use this parameter to display the NVMe I/O queue count, and NVMe I/O queue
depth that match the priority that you specify.

``[-io-queue-count <integer>]`` - I/O Queue count::
Use this parameter to display the priorities, transport-protocol, and NVMe
I/O queue depth on each node that match the NVMe I/O queue count that you
specify.

``[-io-queue-depth <integer>]`` - I/O Queue Depths::
Use this parameter to display the priorities, transport-protocol, and NVMe
I/O count on each node that match the NVMe I/O queue depth that you
specify.

== Examples
```

cluster-1::*> nvme show-host-priority
(vserver nvme show-host-priority)
Node Protocol Priority I/O Queue Count I/O Queue Depth
--- --- --- --- ---
node1 fc-nvme
regular 4 32
high 6 32
nvme-tcp
regular 2 128
high 4 128
node2 fc-nvme
regular 4 32
high 6 32
nvme-tcp
regular 2 128
high 4 128
8 entries were displayed.

```
[[ID94901ac33d83c3723b8cb58ed8d5f03a]]
= vserver nvme show-interface
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display the NVMe over Fabrics LIF configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme show-interface` command displays the currently available
interfaces to the NVMe protocol.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the NVMe interfaces that match the Vserver
that you specify.

``[-lif <lif-name>]`` - Logical Interface::
Use this parameter to display the NVMe interfaces that match the LIF that
you specify.

``[-vserver-uuid <UUID>]`` - Vserver UUID::
Use this parameter to display the NVMe interfaces that match the Vserver
UUID that you specify.

``[-home-node <nodename>]`` - Home Node::
Use this parameter to display the NVMe interfaces that match the home node
that you specify.

``[-home-port {<netport>|<ifgrp>}]`` - Home Port::
Use this parameter to display the NVMe interfaces that match the home port
that you specify.

``[-status-admin {up|down}]`` - Status Admin::
Use this parameter to display the NVMe interfaces that match the
administrative status that you specify.
```

```
``[-physical-protocol {fibre-channel|ethernet}]`` - Physical Protocol::
Use this parameter to display the NVMe interfaces that match the physical
protocol that you specify.

``[-transport-protocol {fc-nvme|nvme-tcp}]`` - Transport Protocol::
Use this parameter to display the NVMe interfaces that match the transport
protocol that you specify.

``[-transport-address <text>]`` - Transport Address::
Use this parameter to display the NVMe interfaces that match the transport
address that you specify.

``[-comment <text>]`` - Comment::
Use this parameter to display the NVMe interfaces that match the textual
comment that you specify.

``[-fc-wwnn <FC WWN>]`` - FC WWNN::
Use this parameter to display the NVMe interfaces that match the FC WWNN
that you specify.

``[-fc-wwpn <FC WWN>]`` - FC WWPN::
Use this parameter to display the NVMe interfaces that match the FC WWPN
that you specify.

``[-lif-id <integer>]`` - LIF ID::
Use this parameter to display the NVMe interfaces that match the LIF ID
that you specify.

``[-lif-uuid <UUID>]`` - LIF UUID::
Use this parameter to display the NVMe interfaces that match the LIF UUID
that you specify.

== Examples
```

cluster1::*> vserver nvme show-interface
Vserver Logical Interface Home Node:Port Transport Protocols
------- ------------------- ----------------- ----------------------------
vs_1
nvme1 node1:1a fc-nvme
Transport Address: nn-0x2000005056b45113:pn-0x2001005056b45113

```
[[ID8bee98b25e552d5a557f4b112b2531dd]]
= vserver nvme show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show NVMe service configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme show` command displays the current status of the NVMe
service in a cluster.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the NVMe services that match the Vserver
that you specify.

``[-status-admin {down|up}]`` - Administrative Status::
Use this parameter to display the NVMe services that match the
administrative status that you specify.

``[-discovery-nqn <text>]`` - Discovery Subsystem NQN::
Use this parameter to display the NVMe services that match the discovery
subsystem NQN that you specify.

``[-mdns-service-discovery-enabled {true|false}]`` - Use Multicast DNS
Service Discovery [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display the NVMe services that match the service
discovery state that you specify.

== Examples
```

cluster1::*> vserver nvme show
Vserver Status Admin
--------- ------------
vs1 up

vs2 up
2 entries were displayed.

```
[[ID780f904f8dba61d1e66f6bcec40d2333]]
= vserver nvme feature show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NVMe target features

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Displays the list of currently enabled NVMe over Fabrics features and
limits.

== Examples
```

cluster1::*> vserver nvme feature show
Enabled Features: nvmf, fc-nvme, multi-path, multi-node
Maximum Namespace Size: 15.97TB
Maximum NSID: 00000200h
Maximum ANAGRPID: 00000200h

```
[[ID3e0471ae6d5a99b682e825c6b3375d4c]]
= vserver nvme namespace convert-from-lun
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Transition an existing LUN into an NVMe namespace.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command transitions an existing LUN in-place into an NVMe namespace
within a volume in a Vserver. When performing an in-place transition, the
application data remains unchanged and only the metadata gets modified.

When using this command, it is necessary that any existing maps to an
Igroup for this LUN be removed first. If any maps exist for the specified
LUN, you will receive an error message.

Only LUNs with a zero prefix and suffix size can be transitioned into an
NVMe namespace. You will receive an error message if the LUN has a non-
zero prefix and/or suffix size.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
The name of the Vserver containing the LUN. If only one data Vserver
exists, you do not need to specify this parameter.

``-lun-path <path>`` - Path of the LUN::
Specifies the path of the LUN you want to transition into an NVMe
namespace. Examples of correct LUN paths: ``_/vol/vol1/lun1_`` and
``_/vol/vol1/qtree1/lun1_``

== Examples
```

cluster1::> vserver nvme namespace convert-from-lun -vserver vs1 -lun-path /vol/vol1/lun1

```
Transitions LUN ``_lun1_`` in-place to NVMe namespace within volume
``_vol1_`` in Vserver ``_vs1_`` .
```

```
[[ID97d97ea579f11e41a98b857c32d83dce]]
= vserver nvme namespace create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an NVMe namespace

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme namespace create` command creates a new NVMe namespace
of a specific size. You must create NVMe namespaces at the root of a
volume or qtree.

When you create an NVMe namespace, it is non-space reserved.

NOTE: This command is not supported for FlexGroup volumes or Vservers with
Infinite Volumes.



== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-path <path>`` - Namespace Path::
Specifies the path of the NVMe namespace. Examples of correct NVMe
namespace paths: ``_/vol/vol1/ns1_`` and ``_/vol/vol1/qtree1/ns1_``

``-size {<integer>[KB|MB|GB|TB|PB]}`` - Size::
Specifies the size of the NVMe namespace in bytes. You can specify a
multiplier suffix:
+
  * KB (1024 bytes)
  * MB (KB*KB bytes)
  * GB (KB*MB bytes)
  * TB (MB*MB bytes)
  * PB (KB*TB bytes)
```

```
``-ostype {aix|linux|vmware|windows}`` - OS Type::
Specifies the operating system type of the NVMe namespace. The OS types
are:
+
  * aix - the NVMe namespace stores AIX file system data.
  * linux - the NVMe namespace stores Linux file system data.
  * vmware - the NVMe namespace stores VMware file system data.
  * windows - the NVMe namespace stores Windows file system data.



``[-comment <text>]`` - Comment::
Contains a textual description of the NVMe namespace.

``[-block-size {512|4KB}]`` - Block Size::
Specifies the block size of the NVMe namespace in bytes. Valid block sizes
are
+
  * ``_512_`` (in ONTAP 9.6 and later)
  * ``_4096_``



``[-foreground {true|false}]`` - Foreground::
 This parameter specifies whether the operation runs in the foreground.
The default setting is true (the operation runs in the foreground). When
set to true, the command will not return until the operation completes.

== Examples
```

cluster1::*> vserver nvme namespace create -vserver vs_1 -path /vol/nsvol/namespace1 -size 100g -ostype linux

```
Creates an NVMe namespace at path ``_/vol/ns/vol/namespace1_`` on Vserver
``_vs_1_`` .




[[IDf2599f8d89dadb78d06de7255c730f68]]
= vserver nvme namespace delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an NVMe namespace

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme namespace delete` command deletes an NVMe namespace from
a specified Vserver and volume.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-path <path>`` - Namespace Path::
Specifies the path of the NVMe namespace. Examples of correct NVMe
namespace paths: ``_/vol/vol1/ns1_`` and ``_/vol/vol1/qtree1/ns1_``

``[-skip-mapped-check <true>]`` - Skip Mapped Check::
This option is required to delete an NVMe namespace that is attached to a
subsystem.

``[-foreground {true|false}]`` - Foreground::
 This parameter specifies whether the operation runs in the foreground.
The default setting is true (the operation runs in the foreground). When
set to true, the command will not return until the operation completes.

== Examples
```

cluster1::> vserver nvme namespace delete -vserver vs1 -path /vol/nsvol/ns1

```
Deletes the NVMe namespace at path ``_/vol/nsvol/ns1_`` on Vserver
``_vs1_`` .




[[ID25e56471f28632f1d998f33bc193bf59]]
= vserver nvme namespace modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an NVMe namespace

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme namespace modify` command modifies NVMe namespace
attributes.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-path <path>`` - Namespace Path::
Specifies the path of the NVMe namespace. Examples of correct NVMe
namespace paths: ``_/vol/vol1/ns1_`` and ``_/vol/vol1/qtree1/ns1_``

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size::
Specifies the size of the NVMe namespace in bytes. You can specify a
multiplier suffix:
+
  * KB (1024 bytes)
  * MB (KB*KB bytes)
  * GB (KB*MB bytes)
  * TB (MB*MB bytes)
  * PB (KB*TB bytes)



``[-comment <text>]`` - Comment::
Contains a textual description of the NVMe namespace.
```

```
``[-state {online|offline|nvfail|space-error}]`` - State [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to modify the state of an NVMe namespace from nvfail to
online. Any other state change is not supported.

``[-skip-shrink-check <true>]`` - Skip checking for NS shrink [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to shrink NVMe namespace size.

== Examples
```

cluster1::*> vserver nvme namespace modify -path /vol/nsvol/ns1 -vserver vs_1 -size 30GB

```
Modifies the size of NVMe namespace at path ``_/vol/nsvol/ns1_`` on
Vserver ``_vs_1_`` to ``_30GB_`` .




[[IDf7ecd9d85cf08355096bcf481e956778]]
= vserver nvme namespace show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NVMe namespaces

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme namespace show` command displays information for NVMe
namespaces.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Selects the NVMe namespaces that match the Vserver that you specify.

``[-path <path>]`` - Namespace Path::
Selects the NVMe namespace that matches the path that you specify.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size::
Selects the NVMe namespaces that match the size that you specify.

``[-size-used {<integer>[KB|MB|GB|TB|PB]}]`` - Size Used::
Selects the NVMe namespaces that match this parameter value.

``[-ostype {aix|linux|vmware|windows}]`` - OS Type::
Selects the NVMe namespaces that match this parameter value.

``[-comment <text>]`` - Comment::
Selects the NVMe namespaces that match this parameter value.

``[-block-size {512|4KB}]`` - Block Size::
Selects the NVMe namespaces that match this parameter value.

``[-state {online|offline|nvfail|space-error}]`` - State::
Selects the NVMe namespaces that match this parameter value.

``[-is-read-only {true|false}]`` - Is Read Only::
Selects the NVMe namespaces that match this parameter value.

``[-creation-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Creation Time::
Selects the NVMe namespaces that match this parameter value.

``[-uuid <UUID>]`` - Namespace UUID::
Selects the NVMe namespaces that match this parameter value.

``[-restore-inaccessible {true|false}]`` - Restore Inaccessible::
Selects the NVMe namespaces that match this parameter value.

``[-node <nodename>]`` - Node Hosting the Namespace::
Selects the NVMe namespaces that match this parameter value.

``[-volume <volume name>]`` - Volume Name::
Selects the NVMe namespaces that match this parameter value.

``[-qtree <qtree name>]`` - Qtree Name::
Selects the NVMe namespaces that match this parameter value.
```

```
``[-subsystem <text>]`` - Mapped Subsystem::
Selects the NVMe namespaces that are attached to a Subsystem that matches
this parameter value.

``[-nsid <Hex 32bit Integer>]`` - Namespace ID::
Selects the NVMe namespaces that match this parameter value.

``[-anagrpid <Hex 32bit Integer>]`` - ANA Group ID [.advaced-
priv]#(privilege: advanced)#::
Selects the NVMe namespaces that match Asymmetric Namespace Access (ANA)
group identifier that you specify.

``[-vserver-id <integer>]`` - Vserver ID::
Selects the NVMe namespaces that match this parameter value.

``[-container-state {online|aggregate-offline|volume-offline|error}]`` -
Namespace Container State [.advaced-priv]#(privilege: advanced)#::
Selects the namespaces that match this parameter value. The container
states are:
+
  * ``_online_`` - The namespace's aggregate and volume are online.
  * ``_aggregate-offline_`` - The namespace's aggregate is offline.
  * ``_volume-offline_`` - The namespace's volume is offline.
  * ``_error_`` - An error occurred accessing the namespace's volume.



``[-include-offline-containers <true>]`` - Include Namespaces on Offline
Volumes and Aggregates [.advaced-priv]#(privilege: advanced)#::
If true, include available information for namespaces in offline
aggregates and offline volumes in the output. By default, namespaces in
offline aggregates and offline volumes are excluded from the output.

``[-application <text>]`` - Application::
Selects the namespaces that are part of an application that matches the
parameter value.

``[-consistency-group <text>]`` - Consistency Group Name::
Selects the namespaces that are associated with the specified consistency
group.

== Examples
```

cluster::*> vserver nvme namespace show -vserver vs1
Vserver Path State Size Subsystem NSID

--------- ---------- ------- ------ --------- ---------
vs1 /vol/test_1_vol/ns1 online 10GB subsys1 00000001h
vs1 /vol/test_1_vol/ns2 online 500MB subsys1 00000002h
vs1 /vol/test_1_vol/ns3 online 1TB - -
3 entries were displayed.

```
[[ID11ee9973a96ba3847401ad1d3f5e6bb4]]
= vserver nvme subsystem create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an NVMe target subsystem

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem create` command creates a new NVMe target
subsystem.

When you create an NVMe subsystem on a Vserver, the Vserver must meet the
following pre-conditions:

* The Vserver must have an NVMe service created.
* The Vserver must not already have an NVMe subsystem by the same name.



NOTE: The NVMe subsystem identifiers are assigned by the system. The NQN
is derived from the Vserver UUID and subsystem name and may not be
specified or modified by the user.



== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-subsystem <text>`` - Subsystem::
Specifies the NVMe target subsystem.

``-ostype {aix|linux|vmware|windows}`` - OS Type::
```

```
Specifies the operating system type of the NVMe subsystem. The OS types
are:
+
  * aix - the initiators belong to an AIX host.
  * linux - the initiators belong to a Linux host.
  * vmware - the initiators belong to a VMware ESX host.
  * windows - the initiators belong to a Windows host.



``[-comment <text>]`` - Comment::
Contains a textual description of the NVMe subsystem.

``[-delete-on-unmap {true|false}]`` - Delete on Unmap [.advaced-
priv]#(privilege: advanced)#::
Specifies whether the subsystem should be deleted automatically when it is
no longer mapped to a namespace.

``[-vendor-uuids <UUID>,...]`` - Vendor UUIDs [.advaced-priv]#(privilege:
advanced)#::
Specifies vendor-specific identifiers (UUIDs) optionally assigned to an
NVMe subsystem when the subsystem is created. The identifiers are used to
enable vendor-specific NVMe protocol features. The identifiers are
provided by a host application vendor and shared with NetApp prior to a
joint product release. Creating an NVMe subsystem with an unknown or non-
specific identifier will have no effect on the NVMe subsystem. Refer to
the ONTAP SAN Administration Guide for a list of the supported vendor-
specific identifiers. After a subsystem is created, the vendor-specific
identifiers cannot be changed or removed.

== Examples
```

cluster1::*> vserver nvme subsystem create -vserver vs_1 -subsystem sub_1 -ostype linux

```
Creates a subsystem named ``_sub_1_`` on Vserver ``_vs_1_`` .




[[ID513aa837c3ddc2fbc7396ddba75ff50]]
= vserver nvme subsystem delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete the subsystem

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem delete` command deletes an NVMe subsystem from
a specified Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-subsystem <text>`` - Subsystem::
Specifies the NVMe target subsystem.

``[-skip-mapped-check <true>]`` - Skip Mapped Namespace Check::
Required to delete an NVMe subsystem with attached NVMe namespaces.

``[-skip-host-check <true>]`` - Skip Host Check::
Required to delete an NVMe subsystem with associated hosts.

== Examples
```

cluster1::*> vserver nvme subsystem delete -vserver vs_1 -subsystem sub_1

```
Deletes the NVMe subsystem ``_sub_1_`` on Vserver ``_vs_1_`` .




[[ID58f5edaee75cdfb688d98c88fbbef718]]
= vserver nvme subsystem modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an NVMe target subsystem

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem modify` command modifies an existing NVMe
target subsystem.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-subsystem <text>`` - Subsystem::
Specifies the NVMe target subsystem.

``[-comment <text>]`` - Comment::
Contains a textual description of the NVMe subsystem.

``[-delete-on-unmap {true|false}]`` - Delete on Unmap [.advaced-
priv]#(privilege: advanced)#::
Specifies whether the subsystem should be deleted automatically when it is
no longer mapped to a namespace.

== Examples
```

cluster1::*> vserver nvme subsystem modify -vserver vs_1 -subsystem sub_1 -comment "Example Comment"

```
Modifies the comment on the subsystem named ``_sub_1_`` on Vserver
``_vs_1_`` .
```

```
[[ID90ce4ff6ac9865988c1dbe369490cb63]]
= vserver nvme subsystem show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NVMe target subsystems

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem show` command displays information for NVMe
subsystems.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the NVMe subsystems that match the Vserver
that you specify.

``[-subsystem <text>]`` - Subsystem::
Use this parameter to display the NVMe subsystem that matches the name
that you specify.

``[-ostype {aix|linux|vmware|windows}]`` - OS Type::
Use this parameter to display the NVMe subsystems that match this
parameter value.

``[-comment <text>]`` - Comment::
Use this parameter to display the NVMe subsystems that match this
```

```
parameter value.

``[-target-nqn <text>]`` - Target NQN::
Use this parameter to display the NVMe subsystems that match this
parameter value.

``[-serial-number <text>]`` - Serial Number::
Use this parameter to display the NVMe subsystems that match this
parameter value.

``[-default-io-queue-count <integer>]`` - Default Number of Host I/O Queue
Pairs::
Specifies the IO queue count inherited by hosts added to this subsystem.
The actual value used when a connection is established may vary depending
on the host and transport protocol used.

``[-default-io-queue-depth <integer>]`` - Default Host I/O Queue Depth::
Specifies the IO queue depth inherited by hosts added to this subsystem.
The actual value used when a connection is established may vary depending
on the host and transport protocol used.

``[-uuid <UUID>]`` - UUID::
Use this parameter to display the NVMe subsystems that have UUIDs matching
the parameter value.

``[-delete-on-unmap {true|false}]`` - Delete on Unmap [.advaced-
priv]#(privilege: advanced)#::
Specifies whether the subsystem should be deleted automatically when it is
no longer mapped to a namespace.

``[-vendor-uuids <UUID>,...]`` - Vendor UUIDs [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display the NVMe subsystems that have vendor-
specific UUIDs that match the parameter value.

== Examples
```

cluster1::*> vserver nvme subsystem show -vserver vs_1
Vserver Subsystem Target NQN
------- ------------ ----------
vs_1
ss1 nqn.1992-08.netapp.com:sn.ccb5a7d5d9d311e7924e005056b45113:subsystem.ss1
ss2 nqn.1992-08.netapp.com:sn.ccb5a7d5d9d311e7924e005056b45113:subsystem.ss2
2 entries were displayed.

```
[[ID4bf47452a09f1de59d55a96705128638]]
```

```
= vserver nvme subsystem controller show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display active NVMe controllers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem controller show` command displays information
for established NVMe controllers. An NVMe controller is established upon
each host connection to a subsystem.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the NVMe controllers that match the Vserver
that you specify.

``[-subsystem <text>]`` - Subsystem::
Use this parameter to display the NVMe controllers that match the
subsystem that you specify.

``[-controller-id <Hex 16bit Integer>]`` - Controller ID::
Use this parameter to display the NVMe controllers that match the
controller ID that you specify. The controller ID is the identifier
assigned by the NVMe subsystem upon host login and controller creation.

``[-lif <text>]`` - Logical Interface::
Use this parameter to display the NVMe controllers that match the LIF that
you specify.

``[-node <nodename>]`` - Node::
```

```
Use this parameter to display the NVMe controllers that match the cluster
node that you specify.

``[-host-nqn <text>]`` - Host NQN::
Use this parameter to display the NVMe controllers that match the host NQN
that you specify. The host NQN is the NVMe identifier assigned to a
specific host.

``[-transport-protocol {fc-nvme|nvme-tcp}]`` - Transport Protocol::
Use this parameter to display the NVMe controllers that match the
transport protocol that you specify.

``[-initiator-transport-address <text>]`` - Initiator Transport Address::
Use this parameter to display the NVMe controllers that match the
initiator transport address that you specify. The initiator transport
address format depends on the transport protocol in use.

``[-host-id <Hex String>]`` - Host Identifier::
Use this parameter to display the NVMe controllers that match the host
identifier that you specify. The host identifier is a 128-bit identifier
assigned to a specific host.

``[-io-queue-count <integer>]`` - Number of I/O Queues::
Use this parameter to display the NVMe controllers that match the I/O
queue count that you specify.

``[-io-queue-depth <integer>,...]`` - I/O Queue Depths::
Use this parameter to display the NVMe controllers that match the I/O
queue depth that you specify.

``[-admin-queue-depth <integer>]`` - Admin Queue Depth::
Use this parameter to display the NVMe controllers that match the
administrative queue depth that you specify.

``[-max-io-size <integer>]`` - Max I/O Size in Bytes::
Use this parameter to display the NVMe controllers that match the maximum
I/O size (in bytes) that you specify.

``[-kato <integer>]`` - Keep-Alive Timeout (msec)::
Use this parameter to display the NVMe controllers that match the keep-
alive timeout value (in msec) that you specify.

``[-subsystem-uuid <UUID>]`` - Subsystem UUID::
Use this parameter to display the NVMe controllers that match the
Subsystem UUID that you specify.

``[-header-digest-enabled {true|false}]`` - Header Digest Enabled::
```

```
Use this parameter to display the NVMe controllers that have header digest
enabled.

``[-data-digest-enabled {true|false}]`` - Data Digest Enabled::
Use this parameter to display the NVMe controllers that have data digest
enabled.

``[-dhchap-hash-function {sha-256|sha-512}]`` - Authentication Hash
Function::
Use this parameter to display the NVMe controllers that connected with the
specified hash function using DH-HMAC-CHAP in-band authentication.

``[-dhchap-dh-group {none|2048-bit|3072-bit|4096-bit|6144-bit|8192-bit}]``
- Authentication Diffie-Hellman Group::
Use this parameter to display the NVMe controllers that connected with the
specified Diffie-Hellman group size using DH-HMAC-CHAP in-band
authentication.

``[-dhchap-mode {none|unidirectional|bidirectional}]`` - Authentication
Mode::
Use this parameter to display the NVMe controllers that connected with the
specified mode for DH-HMAC-CHAP in-band authentication.

``[-trsvcid <text>]`` - Transport Service Identifier::
Use this parameter to display the NVMe controllers that match the
transport service identifier that you specify.

``[-tls-key-type {none|configured}]`` - TLS Key Type::
Use this parameter to display the NVMe controllers that have the specified
NVMe/TLS-TCP encryption key type.

``[-tls-psk-identity <text>]`` - TLS PSK Identity::
Use this parameter to display the NVMe controllers that have the specified
NVMe/TLS-TCP encryption PSK identity.

``[-tls-cipher {TLS-AES-128-GCM-SHA256|TLS-AES-256-GCM-SHA384}]`` - TLS
Cipher::
Use this parameter to display the NVMe controllers that use the specified
NVMe/TLS-TCP encryption cipher.

== Examples
```

cluster1::*> vserver nvme subsystem controller show -vserver vs_1 -subsystem ss1
Vserver Subsystem ID LIF Host NQN
------- --------- ----- ------ ----------
vs_1 ss1
0003h lif1 nqn.2001-08.example.com:nvme:host1

0004h lif1 nqn.2001-08.example.com:nvme:host2
0009h lif2 nqn.2001-08.example.com:nvme:host1
000Ah lif2 nqn.2001-08.example.com:nvme:host2
4 entries were displayed.

```
[[IDe3c28d6114c491ada56ffbb601bc04d5]]
= vserver nvme subsystem host add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a host to a subsystem

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem host add` command adds an FC-NVMe host to an
NVMe subsystem on a Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-subsystem <text>`` - Subsystem::
Specifies the NVMe target subsystem.

``-host-nqn <text>`` - Host NQN::
Specifies the NVMe subsystem host NQN.

``[-priority {regular|high}]`` - Host priority::
Specifies the host priority.

``[-dhchap-host-secret-key <text>]`` - Authentication Host Secret::
Specifies the NVMe DH-HMAC-CHAP in-band authentication host secret.

``[-dhchap-controller-secret-key <text>]`` - Authentication Controller
Secret::
Specifies the NVMe DH-HMAC-CHAP in-band authentication controller secret.

``[-dhchap-hash-function {sha-256|sha-512}]`` - Authentication Hash
Function::
Specifies the NVMe DH-HMAC-CHAP in-band authentication hash function.
```

```
``[-dhchap-dh-group {none|2048-bit|3072-bit|4096-bit|6144-bit|8192-bit}]``
- Authentication Diffie-Hellman Group::
Specifies the NVMe DH-HMAC-CHAP in-band authentication Diffie-Hellman
group size.

``[-tls-configured-psk <text>]`` - TLS Configured PSK::
Specifies the configured pre-shared key (PSK) for the host.

== Examples
```

cluster::*> vserver nvme subsystem host add -vserver vs_1 -subsystem sub_1 -host-nqn nqn.2001-01.com.example:nvme-host1

```
Adds a host with the specified NQN to the NVMe subsystem ``_sub_1_`` on
Vserver ``_vs_1_`` .




[[ID3d30045fa27c7e17dd8a98d72a0744e2]]
= vserver nvme subsystem host remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a host from a subsystem

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem host remove` command removes an FC-NVMe host
from an NVMe subsystem on a Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-subsystem <text>`` - Subsystem::
Specifies the NVMe target subsystem.

``-host-nqn <text>`` - Host NQN::
Specifies the NVMe subsystem host NQN.

== Examples
```

cluster::*> vserver nvme subsystem host remove -vserver vs_1 -subsystem sub_1 -host-nqn nqn.2001-01.com.example:host1

```
Removes the host with the specified NQN from NVMe subsystem ``_sub_1_`` on
Vserver ``_vs_1_`` .
```

```
[[IDde82395fca7a8b6f905d10b7a233e531]]
= vserver nvme subsystem host show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NVMe hosts configured to the subsystem

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem host show` command displays information for
the NVMe subsystem hosts.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the NVMe subsystem hosts that match the
Vserver that you specify.

``[-subsystem <text>]`` - Subsystem::
Use this parameter to display the NVMe subsystem hosts that match the
subsystem that you specify.

``[-host-nqn <text>]`` - Host NQN::
Use this parameter to display the NVMe subsystem host that matches the
subsystem host NQN that you specify.

``[-priority {regular|high}]`` - Host Priority::
Specifies the host priority.

``[-io-queue-count <integer>]`` - Number of I/O Queue Pairs::
Use this parameter to display the NVMe subsystem hosts that match the
```

```
maximum IO queue count that you specify.

``[-io-queue-depth <integer>]`` - I/O Queue Depth::
Use this parameter to display the NVMe subsystem hosts that match the
maximum IO queue depth that you specify.

``[-dhchap-hash-function {sha-256|sha-512}]`` - Authentication Hash
Function::
Specifies the NVMe DH-HMAC-CHAP in-band authentication hash function.

``[-dhchap-dh-group {none|2048-bit|3072-bit|4096-bit|6144-bit|8192-bit}]``
- Authentication Diffie-Hellman Group::
Specifies the NVMe DH-HMAC-CHAP in-band authentication Diffie-Hellman
group size.

``[-dhchap-mode {none|unidirectional|bidirectional}]`` - Authentication
Mode::
Reports the NVMe DH-HMAC-CHAP in-band authentication mode.

``[-tls-key-type {none|configured}]`` - TLS Key Type::
Specifies the type of TLS pre-shared key (PSK) for the host.

== Examples
```

cluster::*> vserver nvme subsystem host show -vserver vs_1 -subsystem sub_1
Vserver Subsystem Host NQN
------- --------- ----------
vs_1 sub_1 nqn.2001-08.com.example:nvme:host1
nqn.2001-08.com.example:nvme:host2
nqn.2001-08.com.example:nvme:host3
3 entries were displayed.

```
[[ID2be424e07a5a0f3ba32ded3f5756079e]]
= vserver nvme subsystem map add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a namespace map

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem map add` command creates an association on an
NVMe namespace to an NVMe subsystem.

When you add an NVMe subsystem map, the following pre-conditions must
hold:

* The NVMe namespace must not be already mapped to a different subsystem.
* There must be an FC-NVMe LIF on the NVMe namespace owning node.



== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-subsystem <text>`` - Subsystem::
Specifies the NVMe target subsystem.

``-path <path>`` - Namespace Path::
Specifies the path of the NVMe namespace. Examples of correct NVMe
namespace paths: ``_/vol/vol1/ns1_`` and ``_/vol/vol1/qtree1/ns1_``

== Examples
```

cluster::*> vserver nvme subsystem map add -vserver vs_1 -subsystem sub_1 -path /vol/nsvol/namespace1

```
Adds an association on the NVMe namespace at the specified path to NVMe
subsystem ``_sub_1_`` on Vserver ``_vs_1_`` .




[[ID5da7f218a8a8001787ee3b306a070dac]]
= vserver nvme subsystem map remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a namespace map

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem map remove` command deletes an association on
an NVMe namespace to an NVMe subsystem.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
Specifies the Vserver.

``-subsystem <text>`` - Subsystem::
Specifies the NVMe target subsystem.

``-path <path>`` - Namespace Path::
Specifies the path of the NVMe namespace. Examples of correct NVMe
namespace paths: ``_/vol/vol1/ns1_`` and ``_/vol/vol1/qtree1/ns1_``

== Examples
```

cluster::*> vserver nvme subsystem map remove -vserver vs_1 -subsystem sub_1 -path /vol/nsvol/namespace1

```
Removes the association on the NVMe namespace at the specified path to
NVMe subsystem ``_sub_1_`` on Vserver ``_vs_1_`` .
```

```
[[IDad1b953d149d76fd00ea3c23e5b54b44]]
= vserver nvme subsystem map show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display namespace maps within the subsystem

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver nvme subsystem map show` command displays information about
NVMe subsystem maps.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
Use this parameter to display the NVMe subsystem maps that match the
Vserver that you specify.

``[-subsystem <text>]`` - Subsystem::
Use this parameter to display the NVMe subsystem maps that match the
subsystem that you specify.

``[-nsid <Hex 32bit Integer>]`` - NSID::
Use this parameter to display the NVMe subsystem maps that match the NVMe
namespace NSID that you specify.

``[-path <path>]`` - Namespace Path::
Use this parameter to display the NVMe subsystem maps that match the NVMe
namespace path that you specify.
```

```
``[-anagrpid <Hex 32bit Integer>]`` - ANA Group ID [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display the NVMe namespaces that match the
Asymmetric Namespace Access (ANA) group identifier that you specify.

``[-namespace-uuid <UUID>]`` - Namespace UUID::
Use this parameter to display the NVMe ubsystem maps that match the NVMe
namespace UUID that you specify.

== Examples
```

cluster-1::*> vserver nvme subsystem map show -vserver vs_1
Vserver Subsystem NSID Namespace Path
---------- --------- --------- ----------
vs_1 sub_1
00000001h /vol/nsvol1/ns1
00000002h /vol/nsvol1/ns2
2 entries were displayed.

```
:leveloffset: -1


= vserver object-store-server commands

:leveloffset: +1


[[IDfe82140ba119688e0a4c3c9f1347a488]]
= vserver object-store-server create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an object store server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server create` command creates an object store
server.

== Parameters
```

```
``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which to create the
object store server. The Vserver must already exist.

``-object-store-server <Object store server name>`` - Object Store Server
Name::
This parameter specifies the name of the object store server. Note that
the object-store-server name must not begin with a bucket name. For
virtual hosted style (VHS) API access, you must use the same hostname as
the server name configured here.

``[-is-http-enabled {true|false}]`` - Accept Connections Over HTTP::
This optional parameter specifies if server should accept HTTP
connections.

``[-is-https-enabled {true|false}]`` - Accept Connections Over HTTPS::
This optional parameter specifies if server should accept HTTPS
connections.

``[-certificate-name <text>]`` - Name of Certificate Used for HTTPS
Connections::
Common name of the certificate used for HTTPS connections.

``[-listener-port <integer>]`` - Object Store Server Listener Port::
Use this parameter to specify the listener port for the object store
server. The default port is ``_80_`` .

``[-secure-listener-port <integer>]`` - Object Store Server Listener Port
for HTTPS::
Use this parameter to specify the secure listener port for the object
store server. The default port is ``_443_`` .

``-status-admin {down|up}`` - Object Store Server Administrative State::
Use this parameter to specify whether the initial administrative status of
the object store server is up or down. The default setting is ``_up_`` .

``[-comment <text>]`` - Object Store Server Description::
This optional parameter specifies a text comment for the object store
server.

``[-default-unix-user <text>]`` - Default UNIX User for NAS Access::
This optional parameter specifies the default UNIX user for name-mapping
from an S3 user to UNIX user during NAS access. The default UNIX user name
is ``_pcuser_`` .

``[-default-win-user <text>]`` - Default Windows User for NAS Access::
This optional parameter specifies the default Windows user for name-
```

```
mapping from an S3 user to Windows user during NAS access.

``[-is-ldap-fastbind-enabled {true|false}]`` - Is LDAP FastBind
Authentication Enabled? [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether LDAP FastBind authentication is enabled
for the object store server.

``[-max-key-time-to-live
{P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W}]`` -
Maximum Value for Key TTL::
Use this parameter to specify the maximum permissible value for object
store user key time-to-live property.

``[-max-lock-retention-period {{<integer> days|years} | none}]`` - Maximum
Object Lock Retention Period::
This parameter specifies the maximum value that can be specified as lock
retention period for an object in an object store bucket.

``[-min-lock-retention-period {{<integer> days|years} | none}]`` - Minimum
Object Lock Retention Period::
This parameter specifies the minimum value that can be specified as lock
retention period for an object in an object store bucket.

== Examples

The following example creates an object store server OSS1 for Vserver vs1.
```

cluster1::> vserver object-store-server create -vserver vs1 -object-store-server OSS1

```
[[ID39eb8a1a98ac08bb8a7b0bc15a1442ae]]
= vserver object-store-server delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an object store server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server delete` command deletes an object store
server.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for the object store
server you want to delete.

== Examples

The following example deletes an object store server for Vserver vs1.
```

cluster1::> vserver object-store-server delete -vserver vs1

```
[[ID46d8eb717269bc270ca55377c1dae1f2]]
= vserver object-store-server modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an object store server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `vserver object-store-server modify` command modifies an object store
server.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for the object store
server which you want to modify.

``[-object-store-server <Object store server name>]`` - Object Store
Server Name::
This parameter specifies the name of the object store server. Note that
the object-store-server name must not begin with a bucket name.

``[-is-http-enabled {true|false}]`` - Accept Connections Over HTTP::
This optional parameter specifies if server should accept HTTP
connections.

``[-is-https-enabled {true|false}]`` - Accept Connections Over HTTPS::
This optional parameter specifies if server should accept HTTPS
connections.

``[-certificate-name <text>]`` - Name of Certificate Used for HTTPS
Connections::
Common name of the certificate used for HTTPS connections.

``[-listener-port <integer>]`` - Object Store Server Listener Port::
This parameter specifies the listener port for the object store server.

``[-secure-listener-port <integer>]`` - Object Store Server Listener Port
for HTTPS::
This parameter specifies the secure listener port for the object store
server.

``[-status-admin {down|up}]`` - Object Store Server Administrative State::
This parameter specifies the administrative status of the object store
server.

``[-comment <text>]`` - Object Store Server Description::
This parameter specifies the text comment for the object store server.

``[-default-unix-user <text>]`` - Default UNIX User for NAS Access::
This optional parameter specifies the default UNIX user used for name-
mapping from an S3 user to UNIX user during NAS access.

``[-default-win-user <text>]`` - Default Windows User for NAS Access::
This optional parameter specifies the default Windows user used for name-
```

```
mapping from an S3 user to Windows user during NAS access.

``[-is-ldap-fastbind-enabled {true|false}]`` - Is LDAP FastBind
Authentication Enabled? [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether LDAP FastBind authentication is enabled
for the object store server.

``[-max-key-time-to-live
{P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W}]`` -
Maximum Value for Key TTL::
This parameter specifies the maximum permissible value for object store
user key time-to-live property.

``[-max-lock-retention-period {{<integer> days|years} | none}]`` - Maximum
Object Lock Retention Period::
This parameter specifies the maximum value that can be specified as lock
retention period for an object in an object store bucket.

``[-min-lock-retention-period {{<integer> days|years} | none}]`` - Minimum
Object Lock Retention Period::
This parameter specifies the minimum value that can be specified as lock
retention period for an object in an object store bucket.

== Examples

The following example modifies the name of the object store server for
Vserver vs1.
```

cluster1::> vserver object-store-server modify -vserver vs1 -object-store-server OSS2

```
[[ID006e872dcb1a144986dd91a243608da8]]
= vserver object-store-server show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display object store servers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server show` command displays information about
```

```
the object store server.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the object store servers for the specified Vserver

``[-object-store-server <Object store server name>]`` - Object Store
Server Name::
If you specify this parameter, the command displays information only for
object store servers that match the specified object store server name.

``[-is-http-enabled {true|false}]`` - Accept Connections Over HTTP::
If you specify this parameter, the command displays information only for
object store servers that accept HTTP connections.

``[-is-https-enabled {true|false}]`` - Accept Connections Over HTTPS::
If you specify this parameter, the command displays information only for
object store servers that accept HTTPS connections.

``[-certificate-name <text>]`` - Name of Certificate Used for HTTPS
Connections::
If you specify this parameter, the command displays information only for
object store servers that match specified certificate name.

``[-listener-port <integer>]`` - Object Store Server Listener Port::
If you specify this parameter, the command displays information only for
object store servers that match the specified listener port.

``[-secure-listener-port <integer>]`` - Object Store Server Listener Port
for HTTPS::
If you specify this parameter, the command displays information only for
object store servers that match the specified secure listener port.

``[-status-admin {down|up}]`` - Object Store Server Administrative State::
If you specify this parameter, the command displays information only for
object store servers that match the specified administrative status.
```

```
``[-comment <text>]`` - Object Store Server Description::
If you specify this parameter, the command displays information only for
object store servers that match the specified comment field.

``[-default-unix-user <text>]`` - Default UNIX User for NAS Access::
If you specify this parameter, the command displays information only for
object store servers that match the specified default UNIX user.

``[-default-win-user <text>]`` - Default Windows User for NAS Access::
If you specify this parameter, the command displays information only for
object store servers that match the specified default Windows user.

``[-is-ldap-fastbind-enabled {true|false}]`` - Is LDAP FastBind
Authentication Enabled? [.advaced-priv]#(privilege: advanced)#::
This parameter specifies whether LDAP FastBind authentication is enabled
for the object store server.

``[-max-key-time-to-live
{P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W}]`` -
Maximum Value for Key TTL::
If you specify this parameter, the command displays information only for
object store servers that match the specified time-to-live value.

``[-max-lock-retention-period {{<integer> days|years} | none}]`` - Maximum
Object Lock Retention Period::
If you specify this parameter, the command displays information only for
object store servers that match the specified maximum lock retention
period.

``[-min-lock-retention-period {{<integer> days|years} | none}]`` - Minimum
Object Lock Retention Period::
If you specify this parameter, the command displays information only for
object store servers that match the specified minimum lock retention
period.

== Examples

The following example displays information of all object store servers:
```

cluster1::> vserver object-store-server show
Vserver: vs3
Object Store Server Name: test.s3.local
Administrative State: up
Listener Port For HTTP: 80
Secure Listener Port For HTTPS: 443
HTTP Enabled: false

HTTPS Enabled: true
Certificate for HTTPS Connections: server_cert
Default UNIX User: pcuser
Default Windows User: win_user
Maximum Object Lock Retention Period: 100 days
Minimum Object Lock Retention Period: 10 days
Comment: Server comment

```
The following example displays information about the object store server
associated with Vserver vs1:
```

cluster1::> vserver object-store-server show -vserver vs1
Vserver: vs1
Object Store Server Name: test.s3.local
Administrative State: up
Listener Port For HTTP: 80
Secure Listener Port For HTTPS: 443
HTTP Enabled: false
HTTPS Enabled: true
Certificate for HTTPS Connections: server_cert
Default UNIX User: pcuser
Default Windows User: win_user
Maximum Object Lock Retention Period: 100 days
Minimum Object Lock Retention Period: 10 days
Comment: Server comment

```
[[IDdb992e3d47e70feba057d5484720ce31]]
= vserver object-store-server audit create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit create` command creates an audit
configuration for a Vserver.

When you create an object store audit configuration, you can also specify
the rotation method. By default, the audit log is rotated based on size.

You can use the time-based rotation parameters in any combination (`-
```

```
rotate-schedule-month` , `-rotate-schedule-dayofweek` , `-rotate-schedule-
day` , `-rotate-schedule-hour` , and `-rotate-schedule-minute` ). The `-
rotate-schedule-minute` parameter is mandatory. All other time-based
rotation parameters are optional.

The rotation schedule is calculated by using all the time-related values.
For example, if you specify only the `-rotate-schedule-minute` parameter,
the audit log files are rotated based on the minutes specified on all days
of the week, during all hours on all months of the year. If you specify
only one or two time-based rotation parameters (say `-rotate-schedule-
month` and `-rotate-schedule-minutes` ), the log files are rotated based
on the minute values that you specified on all days of the week, during
all hours, but only during the specified months. For example, you can
specify that the audit log is to be rotated during the months January,
March, and August on all Mondays, Wednesdays, and Saturdays at 10:30.

If you specify values for both `-rotate-schedule-dayofweek` and `-rotate-
schedule-day` , they are considered independently. For example if you
specify `-rotate-schedule-dayofweek` as Friday and `-rotate-schedule-day`
as 13 then the audit logs would be rotated on every Friday and on the 13th
day of the specified month, not just on every Friday the 13th.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which to create the
audit configuration. The Vserver must already exist.

``-destination <text>`` - Log Destination Path::
This parameter specifies the audit log destination path where consolidated
audit logs are stored. If the path is not valid, the command fails. The
path can be up to 864 characters in length and must have read-write
permissions.

``[-events {data|management}]`` - Categories of Events to Audit::
This parameter specifies the categories of events to be audited. Supported
event categories are: data and management events, The corresponding
parameter values are: ``_data_`` , ``_management_`` .

``[-format <json>]`` - Log Format::
This parameter specifies the output format of the audit logs. By default,
the output format is JSON.

``[-rotate-size {<size>|-}]`` - Log File Size Limit::
This parameter specifies the audit log file size limit. By default, the
audit log is rotated based on size. The default audit log size is 100 MB.
```

```
``[-rotate-schedule-month <cron_month>,...]`` - Log Rotation Schedule:
Month::
This parameter specifies the monthly schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated during
the months January, March, and August, or during all the months. Valid
values are January, February, March, April, May, June, July, August,
September, October, November, December, and all. Specify "all" to rotate
the audit logs every month.

``[-rotate-schedule-dayofweek <cron_dayofweek>,...]`` - Log Rotation
Schedule: Day of Week::
This parameter specifies the daily (day of the week) schedule for rotating
the audit log. For example, you can specify that the audit log is to be
rotated on Tuesdays and Fridays, or during all the days of a week. Valid
values are Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday,
and all. Specify "all" to rotate the audit logs every day.

``[-rotate-schedule-day <cron_dayofmonth>,...]`` - Log Rotation Schedule:
Day::
This parameter specifies the day of the month schedule for rotating the
audit log. For example, you can specify that the audit log is to be
rotated on the 10th and 20th days of a month, or all days of a month.
Valid values range from 1 to 31.

``[-rotate-schedule-hour <cron_hour>,...]`` - Log Rotation Schedule:
Hour::
This parameter specifies the hourly schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at 6 a.m
and 10 a.m. Valid values range from 0 (midnight) to 23 (11:00 p.m.).
Specify "all" to rotate the audit logs every hour.

``[-rotate-schedule-minute <cron_minute>,...]`` - Log Rotation Schedule:
Minute::
This parameter specifies the minute schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at the
30th minute. Valid values range from 0 to 59.

{ ``[-rotate-limit <integer>]`` - Log Files Rotation Limit::
This parameter specifies the audit log files rotation limit. A value of 0
indicates that all the log files are retained. The default value is 0. For
example, if you enter a value of 5, the last five audit logs are retained.

| ``[-retention-duration
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Log Retention
Duration }::
This parameter specifies the audit log files retention duration. A value
```

```
of 0s indicates that all the log files are retained. The default value is
0s. For example, if you enter a value of 5d0h0m, logs more than 5 days old
are deleted.

== Examples

The following examples create an audit configuration for Vserver vs1 using
size-based rotation.
```

cluster1::> vserver object-store-server audit create -vserver vs1 -destination /audit_log -rotate-size 10MB -rotate-limit 5

```
 + +

The following example creates an audit configuration for Vserver vs1 using
time-based rotation. The audit logs are rotated monthly, all days of the
week, at 12:30.
```

cluster1::> vserver object-store-server audit create -vserver vs1 -destination /audit_log -rotate-schedule-month all -rotate-schedule-dayofweek all -rotate-schedule-hour 12 -rotate-schedule-minute 30

```
The following example creates an audit configuration for Vserver vs1 using
time-based rotation. The audit logs are rotated in January, March, May,
July, September, and November on Monday, Wednesday, and Friday, at 6:15,
6:30, 6:45, 12:15, 12:30, 12:45, 18:15, 18:30, and 18:45. The last 6 audit
logs are retained.
```

cluster1::> vserver object-store-server audit create -vserver vs1 -destination /audit_log -rotate-schedule-month January,March,May,July,September,November -rotate-schedule-dayofweek Monday,Wednesday,Friday -rotate -schedule-hour 6,12,18 -rotate-schedule-minute 15,30,45 -rotate-limit 6

```
The following example creates an audit configuration for Vserver vs1 for
auditing object store data access events in the output log format Json.
```

cluster1::> vserver object-store-server audit create -vserver vs1 -destination /audit_log -format json -events data

```
[[IDebe98657a610e128992165f3dbc8b3be]]
= vserver object-store-server audit delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit delete` command deletes the audit
configuration for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver associated with the audit
configuration to be deleted.

``[-force <true>]`` - Force Delete [.advaced-priv]#(privilege:
advanced)#::
This parameter is used to forcibly delete the audit configuration. By
default the setting is `false` .

== Examples

The following example deletes the audit configuration for Vserver vs1.
```

cluster1::> vserver object-store-server audit delete -vserver vs1

```
[[IDd221927c3d3ef1ecda8a285b6356ebf9]]
= vserver object-store-server audit disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable auditing

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit disable` command disables auditing
for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which auditing is to
be disabled. The Vserver audit configuration must already exist.

== Examples

The following example disables auditing for Vserver vs1.
```

cluster1::> vserver object-store-server audit disable -vserver vs1

```
[[ID512da6d3d92afe0ca1572001be878643]]
= vserver object-store-server audit enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable auditing

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit enable` command enables auditing
for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which auditing is to
be enabled. The Vserver audit configuration must already exist.

``[-force <true>]`` - Force Enable [.advaced-priv]#(privilege:
advanced)#::
This parameter is used to ignore errors while enabling auditing.

== Examples

The following example enables auditing for Vserver vs1:
```

cluster1::> vserver object-store-server audit enable -vserver vs1

```
[[IDef744a31739343fb8dae39432b068617]]
= vserver object-store-server audit modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit modify` command modifies an audit
configuration for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which the audit
configuration is to be modified. The Vserver audit configuration must
already exist.
+
If you have configured time-based rotation, modifying one parameter of
time-based rotation schedule does not affect the other parameters. For
example, if the rotation schedule is set to run at Monday 12:30 a.m., and
you modify the `-rotate-schedule-dayofweek` parameter to
Monday,Wednesday,Friday, the new rotation-schedule rotates the audit logs
on Monday, Wednesday, and Friday at 12:30 a.m. To clear time-based
rotation parameters, you must explicitly set that portion to "-". Some
time-based parameters can also be set to "all".

``[-destination <text>]`` - Log Destination Path::
This parameter specifies the audit log destination path where consolidated
audit logs are stored. If the path is not valid, the command fails. The
path can be up to 864 characters in length and must have read-write
permissions.

``[-events {data|management}]`` - Categories of Events to Audit::
This parameter specifies the categories of events to be audited. Supported
event categories are: data and management events. The corresponding
parameter values are: ``_data_`` , ``_management_`` . By default,
``_data_`` events are enabled

``[-format <json>]`` - Log Format::
This parameter specifies the output format of the audit logs. By default,
the output format is JSON.

``[-rotate-size {<size>|-}]`` - Log File Size Limit::
This parameter specifies the audit log file size limit. By default, the
audit log is rotated based on size. The default audit log size is 100 MB.

``[-rotate-schedule-month <cron_month>,...]`` - Log Rotation Schedule:
Month::
This parameter specifies the monthly schedule for rotating the audit log.
```

```
For example, you can specify that the audit log is to be rotated during
the months January, March, and August, or during all the months. Valid
values are January, February, March, April, May, June, July, August,
September, October, November, December, and all. Specify "all" to rotate
the audit logs every month.

``[-rotate-schedule-dayofweek <cron_dayofweek>,...]`` - Log Rotation
Schedule: Day of Week::
This parameter specifies the daily (day of the week) schedule for rotating
the audit log. For example, you can specify that the audit log is to be
rotated on Tuesdays and Fridays, or during all the days of a week. Valid
values are Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday,
and all. Specify "all" to rotate the audit logs every day.

``[-rotate-schedule-day <cron_dayofmonth>,...]`` - Log Rotation Schedule:
Day::
This parameter specifies the day of the month schedule for rotating the
audit log. For example, you can specify that the audit log is to be
rotated on the 10th and 20th days of a month, or all days of a month.
Valid values range from 1 to 31.

``[-rotate-schedule-hour <cron_hour>,...]`` - Log Rotation Schedule:
Hour::
This parameter specifies the hourly schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at 6 a.m
and 10 a.m. Valid values range from 0 (midnight) to 23 (11:00 p.m.).
Specify "all" to rotate the audit logs every hour.

``[-rotate-schedule-minute <cron_minute>,...]`` - Log Rotation Schedule:
Minute::
This parameter specifies the minute schedule for rotating the audit log.
For example, you can specify that the audit log is to be rotated at the
30th minute. Valid values range from 0 to 59.

{ ``[-rotate-limit <integer>]`` - Log Files Rotation Limit::
This parameter specifies the audit log files rotation limit. A value of 0
indicates that all the log files are retained. The default value is 0.

| ``[-retention-duration
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Log Retention
Duration }::
This parameter specifies the audit log files retention duration. A value
of 0s indicates that all the log files are retained. For example, if you
enter a value of 5d0h0m0s, logs more than 5 days old are deleted.

== Examples
```

```
The following example modifies the rotate-size and rotate-limit field for
Vserver vs1.
```

cluster1::> vserver object-store-server audit modify -vserver vs1 -rotate-size 10MB -rotate-limit 3

```
The following example modifies an audit configuration for Vserver vs1
using the time-based rotation method. The audit logs are rotated monthly,
all days of the week, at 12:30.
```

cluster1::> vserver object-store-server audit modify -vserver vs1 -destination /audit_log -rotate-schedule-month all -rotate-schedule-dayofweek all -rotate-schedule-hour 12 -rotate-schedule-minute 30

```
The following example modifies an audit configuration for Vserver vs1 for
auditing object store data events in the output log format Json.
```

cluster1::> vserver object-store-server audit modify -vserver vs1 -format json -events data

```
[[ID504856a0dc5ed72f5cd3dc44882d1a8f]]
= vserver object-store-server audit rotate-log
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Rotate audit log

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit rotate-log` command rotates audit
logs for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which audit logs are
to be rotated. The Vserver audit configuration must already exist.
Auditing must be enabled for the Vserver.

== Examples

The following example rotates audit logs for Vserver vs1.
```

cluster1::> vserver object-store-server audit rotate-log -vserver vs1

```
[[ID0078f1631b8eae1d844d79cbc353e6c8]]
= vserver object-store-server audit show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the audit configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `vserver object-store-server audit show` command displays object store
audit configuration information about Vservers. The command output depends
on the parameter or parameters specified with the command. If you do not
specify any parameters, the command displays the following information
about all the Vservers:

* Vserver name
* Audit state
* Target directory



You can specify the `-fields` parameter to specify which audit
configuration information to display about Vservers. + You can specify
additional parameters to display only information that matches those
parameters. For instance, to display information about the log file
rotation size of a Vserver whose value matches 10 MB, run the command with
the `-rotate-size 10MB` parameter.

You can specify the `-instance` parameter to display audit configuration
information for all Vservers in list form.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-log-save-details  ]``::
You can specify the `-log-save-details` parameter to display the following
information about all the Vservers:
+
  * Vserver name
  * Rotation file size
  * Rotation schedules
  * Rotation limit



| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information about the
specified Vserver.
```

```
``[-state {true|false}]`` - Auditing State::
If you specify this parameter, the command displays information about the
Vservers that use the specified audit state value.

``[-destination <text>]`` - Log Destination Path::
If you specify this parameter, the command displays information about the
Vservers that use the specified destination path.

``[-events {data|management}]`` - Categories of Events to Audit::
If you specify this parameter, the command displays information about the
Vservers that use the specified category of events that are audited. Valid
values are ``_file-ops_`` , ``_cifs-logon-logoff_`` , ``_cap-staging_`` ,
``_file-share_`` , ``_audit-policy-change_`` , ``_user-account_`` ,
``_security-group_`` and ``_authorization-policy-change_`` . ``_audit-
policy-change_`` will appear only in diag mode.

``[-format <json>]`` - Log Format::
If you specify this parameter, the command displays information about the
Vservers that use the specified log format.

``[-rotate-size {<size>|-}]`` - Log File Size Limit::
If you specify this parameter, the command displays information about the
Vservers that use the specified log file rotation size.

``[-rotate-schedule-month <cron_month>,...]`` - Log Rotation Schedule:
Month::
If you specify this parameter, the command displays information about the
Vservers that use the specified month of the time-based log rotation
scheme. Valid values are January, February, March, April, May, June, July,
August, September, October, November, and December.

``[-rotate-schedule-dayofweek <cron_dayofweek>,...]`` - Log Rotation
Schedule: Day of Week::
If you specify this parameter, the command displays information about the
Vservers that use the specified day of the week of the time-based log
rotation scheme. Valid values are Sunday, Monday, Tuesday, Wednesday,
Thursday, Friday, and Saturday.

``[-rotate-schedule-day <cron_dayofmonth>,...]`` - Log Rotation Schedule:
Day::
If you specify this parameter, the command displays information about the
Vservers that use the specified day of the month of the time-based log
rotation scheme. Valid values range from 1 to 31.

``[-rotate-schedule-hour <cron_hour>,...]`` - Log Rotation Schedule:
```

```
Hour::
If you specify this parameter, the command displays information about the
Vservers that use the specified hour of the time-based log rotation
scheme. Valid values range from 0 (midnight) to 23 (11:00 p.m.).

``[-rotate-schedule-minute <cron_minute>,...]`` - Log Rotation Schedule:
Minute::
If you specify this parameter, the command displays information about the
Vservers that use the specified minute of the time-based log rotation
scheme. Valid values range from 0 to 59.

``[-rotate-schedule-description <text>]`` - Rotation Schedules::
If you specify this parameter, the command displays information about the
Vservers that use the specified rotation schedules. This field is derived
from the rotate-time fields.

``[-rotate-limit <integer>]`` - Log Files Rotation Limit::
If you specify this parameter, the command displays information about the
Vservers that use the specified rotation limit value.

``[-retention-duration
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Log Retention
Duration::
If you specify this parameter, the command displays information about the
Vservers audit logs retention duration.

== Examples

The following example displays the name, audit state, event types, log
format, and target directory for all Vservers.
```

cluster1::> vserver object-store-server audit show
Vserver State Event Types Log Format Target Directory
---------- ------ ---------- --------- ----------
vs1 false data json /audit_log

```
The following example displays the Vserver names and details about the
audit log for all Vservers.
```

cluster1::> vserver object-store-server audit show -log-save-details
Rotation Rotation
Vserver File Size Rotation Schedule Limit
---------- -------- ---------- --------
vs1 100MB - 0

```
The following example displays in list form all audit configuration
information about all Vservers.
```

cluster1::> vserver object-store-server audit show -instance
Vserver: vs1
Auditing state: true
Log Destination Path: /audit_log
Categories of Events to Audit: data
Log Format: json
Log File Size Limit: 100MB
Log Rotation Schedule: Month: -
Log Rotation Schedule: Day of Week: -
Log Rotation Schedule: Day: -
Log Rotation Schedule: Hour: -
Log Rotation Schedule: Minute: -
Rotation Schedules: -
Log Files Rotation Limit: 0
Log Retention Time: 0s

```
[[IDd599ada859c601376185509bc435840b]]
= vserver object-store-server audit event-selector create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an object store server audit event-selector

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit event-selector create` command
creates an audit event-selector for the object store server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the bucket audit
event-selector needs to be created for the object store server bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the audit event-selector needs to be created. The object store
bucket must already exist.

``-access {read-only|write-only|all}`` - Access type for audit events::
Use this parameter to specify which type of event access is being audited.
Possible values are: read-only, write-only or all.

``-permission {allow-only|deny-only|all}`` - Permission type for audit
events::
Use this parameter to specify which type of event permission is being
audited. Possible value are: allow-only, deny-olnly or all.

== Examples

The following example displays information on object store server audit
event-selector for vserver vs1 and bucket bucket1:
```

cluster1::> vserver object-store-server audit event-selector create
-vserver vs1 -bucket bucket1 -access read-only -permission allow-only

```
[[ID52f72c52d0c92b2139067865d2620b43]]
= vserver object-store-server audit event-selector delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an object store server audit event-selector

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit event-selector delete` command
delete an audit event-selector for the object store server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the bucket audit
event-selector needs to be deleted.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the audit event-selector needs to be deleted.

== Examples

The following example delete an object store server audit event-selector
for Vserver vs1 and bucket1:
```

cluster1::> vserver object-store-server audit event-selector delete
-vserver vs1 -bucket bucket1

```
[[ID59042bc6d8dfff264c09b9ef72629bb6]]
= vserver object-store-server audit event-selector modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an object store server audit event-selector

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit event-selector modify` command
modifies an audit event-selector for the object store server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the bucket audit
event-selector needs to be modified for the object store server bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the audit event-selector needs to be modified.

``[-access {read-only|write-only|all}]`` - Access type for audit events::
Use this parameter to specify which type of event access is being audited.
Possible values are: read-only, write-only or all.

``[-permission {allow-only|deny-only|all}]`` - Permission type for audit
events::
Use this parameter to specify which type of event permission is being
audited. Possible value are: allow-only, deny-olnly or all.

== Examples

The following example modified an object store server audit event-selector
for Vserver vs1 and bucket1 with read-only access to write-only access:
```

cluster1::> vserver object-store-server audit event-selector modify
-vserver vs1 -bucket bucket1 -access write-only

```
[[IDded2fc7f70c80f664819335b3575f164]]
= vserver object-store-server audit event-selector show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display object store server audit event-selector

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server audit event-selector show` command
displays information about object store server audit event-selector.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information on the
object store server audit event-selector for the specified Vserver.

``[-bucket <TextNoCase>]`` - Object Store Server Bucket Name::
If you specify this parameter, the command displays information on the
object store server audit event-selector for the specified bucket.

``[-access {read-only|write-only|all}]`` - Access type for audit events::
If you specify this parameter, the command displays information on the
object store server audit event-selector that match the specified access.

``[-permission {allow-only|deny-only|all}]`` - Permission type for audit
events::
If you specify this parameter, the command displays information on the
object store server audit event-selector that match the specified
permission.
```

```
== Examples

The following example displays information on object store server audit
event-selector for vserver vs1 and bucket bucket1:
```

cluster1::> vserver object-store-server audit event-selector show
-vserver vs1 -bucket bucket1
Vserver Bucket Access Permission
---------- --------- ---------- ---------
vs1
bucket1 read-only allow-only

```
The following example displays detailed information of the object server
audit event-selector associated with Vserver vs1.
```

cluster1::> vserver object-store-server audit event-selector show
-vserver vs1

Vserver :vs1
Bucket :bucket1
Access :all
Permission :all

```
[[IDf97aef902b234a5b012ea0fa7b509cb2]]
= vserver object-store-server bucket clear-snapshot-metadata
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all snapshot metadata for the S3 bucket.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver object-store-server bucket clear-snapshot-metadata` command
removes all snapshot metadata for a specified S3 bucket. This command is
available in advanced privilege.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the Vserver that contains the S3 bucket for
which the snapshot metadata will be removed.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the name of the S3 bucket for which the snapshot metadata
will be removed.

== Examples

The following example removes all snapshot metadata for an S3 bucket
bucket-1 on Vserver named vs0.
```

cluster1::> vserver object-store-server bucket clear-snapshot-metadata -vserver vs0 -bucket bucket-1

```
[[ID00fe38ccaf7ddbdf31169b3901ef1b4e]]
= vserver object-store-server bucket create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Create an object store server bucket

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket create` command creates a bucket
for the object store server.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server where the bucket needs to be created. The object store server must
already exist.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket. Note
that the bucket name must not be same as the beginning of the object-
store-server name present in the vserver.

``[-type {s3|nas}]`` - Type of bucket::
This parameter specifies the type of the bucket. The default value is
``_s3_`` .

``[-versioning-state {disabled|enabled|suspended}]`` - Object Store Server
Versioning State::
Use this parameter to specify the state of versioning on the bucket.

``[-comment <text>]`` - Object Store Server Bucket Comment::
This optional parameter specifies a text comment for the object store
server bucket.

{ ``[-aggr-list <aggregate name>,...]`` - List of Aggregates for FlexGroup
Constituents [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the list of aggregates for the FlexGroup
constituents on which the bucket needs to be created. Each entry in the
list will create a constituent on the specified aggregate. The root
constituent will always be placed on the first aggregate in the list,
unless `optimize-aggr-list` is specified as ``_true_`` . An aggregate may
be specified multiple times to have multiple constituents created on it.
This parameter only applies to FlexGroup volumes. This parameter is
available only in advanced privilege mode.

``[-aggr-list-multiplier <integer>]`` - Aggregate List Repeat Count
```

```
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the number of FlexGroup constituents to be
created. This parameter is available only in advanced privilege mode.

``[-optimize-aggr-list {true|false}]`` - Have the System Optimize the
Order of the Aggregate List [.advaced-priv]#(privilege: advanced)#::
Specifies whether to create the constituents of the FlexGroup volume on
which the bucket needs to be created, on the aggregates specified in the
`aggr-list` in the order they are specified, or whether the system should
optimize the ordering of the aggregates. If this value is ``_true_`` , the
system will optimize the ordering of the aggregates specified in the
`aggr-list` . If this value is ``_false_`` the order of the `aggr-list`
will be unchanged. The default value is ``_false_`` . This parameter only
applies to FlexGroup volumes. This parameter is available only in advanced
privilege mode.

{ ``[-used-as-capacity-tier {true|false}]`` - Is Used as Capacity Tier::
Use this parameter to specify if the bucket is going to be used for
capacity tier.

| ``[-storage-service-level <text>]`` - Storage Service Level of the
Bucket }::
Use this parameter to specify the storage service level with which the
bucket should be created.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of the Bucket::
Use this parameter to specify the size of the FlexGroup volume to be
created.

``[-exclude-aggr-list <aggregate name>,...]`` - List of Aggregates to
Exclude During FlexGroup Create::
Use this parameter to specify the list of aggregates to exclude during
FlexGroup creation. This parameter is used only when creating bucket for
capacity tier use case within the local cluster.

``[-qos-policy-group <text>]`` - QoS policy group::
A policy group defines measurable service level objectives (SLOs) that
apply to the storage objects with which the policy group is associated. If
you do not assign a policy group to a bucket, the system wil not monitor
and control the traffic to it.

``[-nas-path <text>]`` - NAS Path corresponding to the Bucket::
This parameter specifies the path to the NAS directory which the bucket
maps to.

``[-retention-mode {no-lock|compliance|governance}]`` - Bucket Retention
```

```
Mode::
Use this parameter to specify the retention-mode in which objects within
the bucket can be locked.

``[-use-mirrored-aggregates {true|false}]`` - Use Mirrored Aggregates::
Use this parameter to specify whether mirrored aggregates are selected for
the FlexGroup on which the bucket will be created. Only mirrored
aggregates are used if this parameter is set to ``_true_`` and only
unmirrored aggregates are used if this parameter is set to ``_false_`` .
The default value is ``_true_`` for a MetroCluster configuration and is
``_false_`` for a non-MetroCluster configuration.

``[-default-retention-period {{<integer> days|years} | none}]`` - Bucket
Default Retention Period::
Use this parameter to specify the retention-period to be applied on all
unlocked objects inserted into the bucket. The retention period can be in
years, or days. A period specified for years and days is represented in
the ISO-8601 format as "10 years" and "100 days" respectively, for example
"10 years" represents a duration of 10 years. The period string must
contain only a single time element that is, either years, or days. A
duration which combines different periods is not supported, for example
"10 years 12 days" is not supported.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
This parameter specifies the snapshot policy for the object store server
bucket.

== Examples

The following example creates an object store server bucket for Vserver
vs1 of size 1TB.
```

cluster1::> vserver object-store-server bucket create -vserver vs1 -bucket testbucket -size 1TB.

```
The following example creates an object store server bucket for Vserver
vs1 of size 1TB using aggr-list.
```

cluster1::*> vserver object-store-server bucket create -vserver vs1 -bucket testbucket -aggr-list aggr1 -size 1TB.

```
[[ID39725ab4bd33a8e3a870d90e029d5ec2]]
= vserver object-store-server bucket delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an object store bucket

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket delete` command deletes the bucket
belonging to the object store server.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server's bucket you want to delete.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the bucket of the object store server
you want to delete.

== Examples

The following example deletes an object store server bucket for Vserver
vs1.
```

cluster1::> vserver object-store-server delete -vserver vs1 -bucket testbucket

```
[[ID7e6a18161722f3b422fb2af98034d687]]
= vserver object-store-server bucket evict-remote-cached-objects
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Evict remote read-write cached objects

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver object-store-server evict-remote-cached-objects` command will
evict read-write dirty cached objects in all constituent volumes of a
given object store server bucket. This command will evict only objects
that are cached on a different volume than its origin volume. This command
requires two parameters - a Vserver name and an object store server bucket
name.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the Vserver for the object store server.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name [.advaced-
priv]#(privilege: advanced)#::
This specifies the name of the object store server bucket.

== Examples

The following example starts the command:
```

cluster1::>vserver object-store-server bucket evict-remote-cached-objects -vserver my-vserver -bucket my-bkt

```
[[ID535a37e2cf95edf8e0ef9431219e70d1]]
= vserver object-store-server bucket modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Modify an object store server bucket

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket modify` command modifies an object
store server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server which you want to modify.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket which
you want to modify.

``[-versioning-state {disabled|enabled|suspended}]`` - Object Store Server
Versioning State::
Use this parameter to specify the state of versioning on the bucket. Note
that the versioning state cannot be modified to 'disabled' from any other
state.

``[-comment <text>]`` - Object Store Server Bucket Comment::
This parameter specifies the text comment for the object store server
bucket.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of the Bucket::
This parameter specifies the size of the object store server bucket.

``[-qos-policy-group <text>]`` - QoS policy group::
A policy group defines measurable service level objectives (SLOs) that
apply to the storage objects with which the policy group is associated. If
you do not assign a policy group to a bucket, the system wil not monitor
and control the traffic to it.

``[-nas-path <text>]`` - NAS Path corresponding to the Bucket::
This parameter specifies the path to the NAS directory which the bucket
maps to.

``[-default-retention-period {{<integer> days|years} | none}]`` - Bucket
Default Retention Period::
Use this parameter to modify the default-retention-period to be applied on
```

```
all unlocked objects to be inserted into the bucket.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
This parameter specifies the snapshot policy for the object store server
bucket.

== Examples

The following example modifies the comment of the object store server
bucket for Vserver vs1.
```

cluster1::> vserver object-store-server bucket modify -vserver vs1 -bucket testbucket -comment test

```
[[ID8e79750e00f602ae710123a8b7067581]]
= vserver object-store-server bucket show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display object store server buckets

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket show` command displays information
about the object store server bucket.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information only about
the object store server buckets for the specified Vserver
```

```
``[-bucket <TextNoCase>]`` - Object Store Server Bucket Name::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified bucket.

``[-type {s3|nas}]`` - Type of bucket::
This parameter specifies the type of bucket.

``[-versioning-state {disabled|enabled|suspended}]`` - Object Store Server
Versioning State::
This parameter specifies the state of the versioning on a bucket.

``[-uuid <UUID>]`` - Object Store Server Bucket UUID::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified bucket uuid.

``[-comment <text>]`` - Object Store Server Bucket Comment::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified comment.

``[-volume <volume name>]`` - Hosting Volume Name::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified volume.

``[-size {<integer>[KB|MB|GB|TB|PB]}]`` - Size of the Bucket::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified size.

``[-logical-used {<integer>[KB|MB|GB|TB|PB]}]`` - Object Store Server
Bucket Logical Used Size::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified logical used size.

``[-object-count <integer>]`` - Object Store Server Object Count::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified object-count.

``[-incomplete-multipart-uploads-count <integer>]`` - Count of Incomplete
Multipart Uploads::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified incomplete-multipart-
uploads-count.

``[-encryption {true|false}]`` - Is Encryption Enabled on Bucket::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified encryption field.

``[-qos-policy-group <text>]`` - QoS policy group::
```

```
A policy group defines measurable service level objectives (SLOs) that
apply to the storage objects with which the policy group is associated. If
you do not assign a policy group to a bucket, the system wil not monitor
and control the traffic to it.

``[-is-protected {true|false}]`` - Is bucket a FabricLink source and
protected::
This parameter specifies whether a bucket is protected using SnapMirror
relationship to another bucket.

``[-is-protected-on-ontap {true|false}]`` - Is bucket protected over
ONTAP::
This parameter specifies whether a bucket is protected using SnapMirror
relationship to another ONTAP bucket.

``[-is-protected-on-cloud {true|false}]`` - Is bucket protected over
Cloud::
This parameter specifies whether a bucket is protected using SnapMirror
relationship to a Cloud bucket.

``[-is-protected-on-external-cloud {true|false}]`` - Is bucket protected
on External Cloud::
This parameter specifies whether a bucket is protected using S3 SnapMirror
relationship to a bucket on an external Cloud provider i.e. excluding
providers types ONTAP_S3 and SGWS.

``[-nas-path <text>]`` - NAS Path corresponding to the Bucket::
This parameter specifies the path to the NAS directory which the bucket
maps to.

``[-retention-mode {no-lock|compliance|governance}]`` - Bucket Retention
Mode::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified retention-mode. This
parameter specifies the object locking mode of the bucket.

``[-default-retention-period {{<integer> days|years} | none}]`` - Bucket
Default Retention Period::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified default-retention-
period field. This parameter specifies the default-retention-period
applied on the bucket.

``[-snapshot-policy <snapshot policy>]`` - Snapshot Policy::
If you specify this parameter, the command displays information only for
object store server buckets that match the specified snapshot-policy
```

```
field. This parameter specifies the snapshot-policy for the bucket.

== Examples

The following example displays information of all object store servers
buckets:
```

cluster1::> vserver object-store-server bucket show
nsankaracluster-1::*> vserver object-store-server bucket show
Vserver Bucket Type Volume Size Encryption Role NAS Path
---------- --------- ------- ---------- ---------- ---------- ---------- ---------
vs1 s3bucket1 nas s3adapter - false - /s3adapter
vs1 testbucket1 s3 fg_oss_1654817100 1.56GB false standalone -
Comment: test1
vs2 nasbucket1 nas vol2 - false - /vol2
vs2 nasbucket2 nas vol2 - false - /vol2
4 entries were displayed.

```
The following example displays information of the object store server
bucket associated with Vserver vs1:
```

cluster1::> vserver object-store-server bucket show -vserver vs1
Vserver Bucket Type Volume Size Encryption Role NAS Path
---------- --------- ------- ---------- ---------- ---------- ---------- ---------
vs1 s3bucket1 nas s3adapter - false - /s3adapter
vs1 testbucket1 s3 fg_oss_1654817100 1.56GB false standalone -
Comment: test1

```
The following example displays detailed information of the object store
server bucket associated with Vserver vs1:
```

cluster1::> vserver object-store-server bucket show -vserver vs1 -instance
Vserver Name: vs1
Object Store Server Bucket Name: s3bucket1
Type of bucket: nas
Object Store Server Versioning State: -
Object Store Server Bucket UUID: c621d53a-ddf0-11ec-958f-005056bba281
Object Store Server Bucket Comment:
Hosting FlexGroup Volume Name: s3adapter
Size of the Bucket: -
Object Store Server Bucket Logical Used Size: -
Object Store Server Object Count: -
Is Encryption Enabled on Bucket: false
QoS policy group: -
Role of the Bucket: -
Is bucket a FabricLink source and protected: -
Is bucket protected over ONTAP: -
Is bucket protected over Cloud: -

Is bucket protected on External Cloud: -
NAS Path corresponding to the Bucket: /s3adapter
Vserver Name: vs1
Object Store Server Bucket Name: testbucket1
Type of bucket: s3
Object Store Server Versioning State: disabled
Object Store Server Bucket UUID: 6a1fa354-e84b-11ec-adce-005056bba281
Object Store Server Bucket Comment:
Hosting FlexGroup Volume Name: fg_oss_1654817100
Size of the Bucket: 1.56GB
Object Store Server Bucket Logical Used Size: 0B
Object Store Server Object Count: 0
Is Encryption Enabled on Bucket: false
QoS policy group: -
Role of the Bucket: standalone
Is bucket a FabricLink source and protected: false
Is bucket protected over ONTAP: false
Is bucket protected over Cloud: false
Is bucket protected on External Cloud: false
NAS Path corresponding to the Bucket: -
2 entries were displayed.

```
[[ID04e73b2fce22e6737388c7676579fac8]]
= vserver object-store-server bucket cors-rule create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a Cross-Origin rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket cors-rule create` command creates
a cross-origin resource sharing (CORS) rule for the object store server
bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver hosting the object store
server bucket for which the CORS rule will to be created.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
```

```
which the CORS rule will to be created.

``[-index <integer>]`` - Cross-Origin Access Rule Index::
This parameter specifies the index of the object store server bucket CORS
rule. This is an optional parameter. Its value is calculated automatically
by the system, if not specified.

``[-rule-id <text>]`` - Cross-Origin Access Rule ID::
This parameter specifies the ID of the object store server bucket CORS
rule.

``-allowed-origins <text>,...`` - List of the Origins Allowed for Cross-
Origin Requests::
This parameter specifies a list of origins from where a cross-origin
request is allowed to originate from for the object store bucket.

``-allowed-methods <text>,...`` - List of HTTP Request Methods Allowed in
Cross-Origin Requests::
This parameter specifies a list of HTTP methods allowed in the cross-
origin requests.

``[-allowed-headers <text>,...]`` - List of the Headers Allowed in Cross-
Origin Requests::
This parameter specifies a list of HTTP headers allowed in the cross-
origin requests.

``[-expose-headers <text>,...]`` - Headers that Servers Allow Clients to
Access::
This parameter specifies a list of extra headers sent in the response that
customers can access from their applications.

``[-max-age-in-seconds <integer>]`` - How Long ( in Seconds ) the Results
of a Cross-Origin Pre-Flight Request Can Be Cached?::
This optional parameter specifies the time in seconds for your browser to
cache the pre-flight response for the specified resource.

== Examples

The following example creates an object store server bucket CORS rule for
Vserver vs1 and bucket bb1.
```

cluster1::> vserver object-store-server bucket cors-rule create -vserver vs1 -bucket bb1 -allowed-origins example.com -allowed-methods POST,DELETE -allowed-headers x-amz-request-id -max-age-in-seconds 2000 -expose-headers x-purpose -rule-id rule1

```
[[ID0156f817eae79b1392d43d76d0c9cb2f]]
= vserver object-store-server bucket cors-rule delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Cross-Origin rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket cors-rule delete` command deletes
a cross-origin resource sharing (CORS) rule belonging to the object store
server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver hosting the object store
server bucket for which the CORS rule will be deleted.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the CORS rule will be deleted.

``-index <integer>`` - Cross-Origin Access Rule Index::
This parameter specifies the index of the object store server bucket CORS
rule.

== Examples

The following example deletes an object store server bucket CORS rule with
index 1 of Vserver vs1 and bucket bucket1.
```

cluster1::> vserver object-store-server bucket cors-rule delete -vserver vs1 -bucket bucket1 -index 1

```
[[ID79582a79db4cd1c6ea77a0b0a5529b0c]]
= vserver object-store-server bucket cors-rule show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the Cross-Origin rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket cors-rule show` command displays
information about the cross-origin resource sharing (CORS) rules for the
specified object store server bucket.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules for the specified Vserver.

``[-bucket <TextNoCase>]`` - Object Store Server Bucket Name::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules for the specified object store
server bucket.

``[-index <integer>]`` - Cross-Origin Access Rule Index::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules that match the specified index.

``[-rule-id <text>]`` - Cross-Origin Access Rule ID::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules that match the specified rule ID.

``[-allowed-origins <text>,...]`` - List of the Origins Allowed for Cross-
Origin Requests::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules that match the specified allowed
origins.
```

```
``[-allowed-methods <text>,...]`` - List of HTTP Request Methods Allowed
in Cross-Origin Requests::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules that match the specified allowed
methods.

``[-allowed-headers <text>,...]`` - List of the Headers Allowed in Cross-
Origin Requests::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules that match the specified allowed
headers.

``[-expose-headers <text>,...]`` - Headers that Servers Allow Clients to
Access::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules that match the specified exposed
headers.

``[-max-age-in-seconds <integer>]`` - How Long ( in Seconds ) the Results
of a Cross-Origin Pre-Flight Request Can Be Cached?::
If you specify this parameter, the command displays information on the
object store server bucket CORS rules that match the specified value for
this parameter.

== Examples

The following example displays information on object store server bucket
CORS rules for vserver vs1 and bucket bucket1:
```

cluster1::> vserver object-store-server bucket cors-rule show -vserver vs1 -bucket bb1
Vserver Bucket Rule-Index Origin Methods
-------- --------- --------------- ---------------- ---------------
vs1 bb1 1 hello.com POST
vs1 bb1 2 * POST
vs1 bb1 5 example.com POST, DELETE
2 entries were displayed.

```
The following example displays detailed information of the object store
server bucket CORS rule associated with Vserver vs1 and bucket bucket1:
```

cluster1::> vserver object-store-server bucket cors-rule show -vserver vs1 -bucket bb1 -index 5

Vserver: vs1
Object Store Server Bucket Name: bb1
Cross-Origin Access Rule Index: 5
Cross-Origin Access Allowed Origins: example.com

Cross-Origin Access Allowed Methods: POST, DELETE
Cross-Origin Access Allowed Headers: x-amz-request-id
Cross-Origin Access Rule Identifier: rule1
Cross-Origin Access Exposed Headers: x-purpose
Time a browser can cache Cross-Origin Response: 2000

```
[[IDcb1612ab0929bd1acc04d78ea2006e06]]
= vserver object-store-server bucket lifecycle-management-rule create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a lifecycle management rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket lifecycle-management-rule create`
command creates a lifecycle management rule for the object store server
bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the bucket
lifecycle management rule needs to be created for the object store server
bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the lifecycle management rule needs to be created. The object store
server bucket must already exist.

``-rule-id <text>`` - Lifecycle Management Rule Identifier::
This parameter specifies the rule identifier of the lifecycle management
rule to be applied on the object store server bucket.

``[-index <integer>]`` - Lifecycle Management Rule Index::
This parameter specifies the index of the lifecycle management rule to be
applied on the object store server bucket.

``[-is-enabled {true|false}]`` - Is This Rule Enabled?::
This parameter specifies whether the configured lifecycle management rule
```

```
is enabled or disabled on the object store server bucket. If you do not
specify this parameter, the default is ``_true_`` .

``[-prefix <text>]`` - Prefix to be Matched with Object Names::
Use this parameter to specify a prefix that is matched against object-
names within a bucket.

``[-tags <text>,...]`` - Tags in Format <tag> or <tag=value>::
Use this parameter to specify a list of key-value paired tags.

``[-obj-size-greater-than {<integer>[KB|MB|GB|TB|PB]}]`` - Min Size of the
Object::
Use this parameter to specify the minimum size of the object for which the
corresponding lifecycle rule is to be applied.

``[-obj-size-less-than {<integer>[KB|MB|GB|TB|PB]}]`` - Max Size of the
Object::
Use this parameter to specify the maximum size of the object for which the
corresponding lifecycle rule is to be applied.

``-action
{Expiration|NoncurrentVersionExpiration|AbortIncompleteMultipartUpload|Exp
irationInternal}`` - Lifecycle Management Action::
Use this parameter to specify lifecycle management actions. The set of
actions that the object store server supports are ``_Expiration_`` ,
``_NoncurrentVersionExpiration_`` and ``_AbortIncompleteMultipartUpload_``
.

{ ``[-obj-age-days <integer>]`` - Number of Days since Creation, After
Which Current Version of Objects Can be Deleted::
Minimum lifetime in number of days since creation, after which objects can
be deleted. This parameter is available for expiration actions only.

``[-obj-exp-date <MM/DD/YYYY HH:MM:SS>]`` - Specific Date When the Objects
Should Expire::
Expiration date of an object. This parameter is available for expiration
actions only.

``[-expired-obj-del-marker {true|false}]`` - Cleanup Object Delete
Markers::
When set to ``_true_`` , an object with a delete marker will be deleted.
This parameter is available for expiration actions only.

| ``[-new-non-curr-versions <integer>]`` - Number of Latest Non-current
Versions to Be Retained::
This parameter specifies the number of latest non-current versions to be
retained. This parameter is available for non-current version expiration
```

```
actions only.

``[-non-curr-days <integer>]`` - Number of Days after Which Non-current
Versions will Be Deleted::
This parameter specifies the number of days after which non-current
versions can be deleted. This parameter is available for non-current
version expiration actions only.

| ``[-after-initiation-days <integer>]`` - Number of Days of Initiation,
After Which Upload Can Be Aborted }::
This parameter specifies the number of days of initiation, after which
uploads can be aborted. This parameter is required for abort-incomplete
multipart upload actions only.

== Examples

The following example creates an object store server bucket lifecycle
management rule for Vserver vs1 and bucket1 which specifies an expiration
action on a set of objects.
```

cluster1::> vserver object-store-server bucket lifecycle-management-rule create -vserver vs1 -bucket bucket1 -rule-id rule1 -prefix obj1/ -action Expiration -obj-age-days 100"

```
[[IDbc6b6925d5492f1441ea60b3b9de2e33]]
= vserver object-store-server bucket lifecycle-management-rule delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Lifecycle Management rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket lifecycle-management-rule delete`
command deletes a lifecycle management rule for the object store server
bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the bucket
```

```
lifecycle management rule needs to be deleted for the object store server
bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the lifecycle management rule needs to be deleted. The object store
server bucket must already exist.

``-rule-id <text>`` - Lifecycle Management Rule Identifier::
This parameter specifies the rule identifier of the lifecycle management
rule to be deleted from the object store server bucket.

``-index <integer>`` - Lifecycle Management Rule Index::
This parameter specifies the index of the lifecycle management rule to be
deleted from the object store server bucket.

``[-force <true>]`` - Ignore Errors::
If this parameter is specified and set to true, the user is not prompted
to confirm each deletion operation. In addition, several potential errors
are ignored. By default, this setting is ``_true_`` .

== Examples

The following example deletes an object store server bucket lifecycle
management rule for Vserver vs1 and bucket1.
```

cluster1::> vserver object-store-server bucket lifecycle-management-rule delete -vserver vs1 -bucket bucket1 -rule-id rule1 -index 1"

```
[[ID15507de7b25e162911bc85e3ec8de102]]
= vserver object-store-server bucket lifecycle-management-rule modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a lifecycle management rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket lifecycle-management-rule modify`
command modifies a lifecycle management rule for the object store server
```

```
bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the bucket
lifecycle management rule will be modified for the object store server
bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the lifecycle management rule will be modified. The object store
server bucket must already exist.

``-rule-id <text>`` - Lifecycle Management Rule Identifier::
This parameter specifies the rule identifier of the lifecycle management
rule to be applied on the object store server bucket.

``-index <integer>`` - Lifecycle Management Rule Index::
This parameter specifies the index of the lifecycle management rule to be
applied on the object store server bucket.

``[-is-enabled {true|false}]`` - Is This Rule Enabled?::
This parameter specifies whether the configured lifecycle management rule
is to be enabled or disabled on the object store server bucket. The
default value when a rule is created is ``_true_`` .

{ ``[-obj-age-days <integer>]`` - Number of Days since Creation, After
Which Current Version of Objects Can be Deleted::
This parameter specifies the minimum lifetime, in days since creation,
after which objects can be deleted. This parameter is available for
expiration actions only.

``[-obj-exp-date <MM/DD/YYYY HH:MM:SS>]`` - Specific Date When the Objects
Should Expire::
This parameter specifies the expiration date when an object will expire.
This parameter is available for expiration actions only.

``[-expired-obj-del-marker {true|false}]`` - Cleanup Object Delete
Markers::
This parameter specifies whether to delete an object that has a delete
marker or not. When set to ``_true_`` , an object with a delete marker
will be deleted. This parameter is available for expiration actions only.

| ``[-new-non-curr-versions <integer>]`` - Number of Latest Non-current
Versions to Be Retained::
This parameter specifies the number of latest non-current versions to be
```

```
retained. This parameter is available for non-current version expiration
actions only.

``[-non-curr-days <integer>]`` - Number of Days after Which Non-current
Versions will Be Deleted::
This parameter specifies the number of days after which non-current
versions can be deleted. This parameter is available for non-current
version expiration actions only.

| ``[-after-initiation-days <integer>]`` - Number of Days of Initiation,
After Which Upload Can Be Aborted }::
This parameter specifies the number of days of initiation, after which
uploads can be aborted. This parameter is available for abort-incomplete
multipart upload actions only.

== Examples

The following example modifies an object store server bucket lifecycle
management rule for Vserver vs1 and bucket1 which specifies expiration
action on a set of objects.
```

cluster1::> vserver object-store-server bucket lifecycle-management-rule modify -vserver vs1 -bucket bucket1 -rule-id rule1 -index 1 -obj-age-days 200"

```
[[ID3341804120f8c85211562746796e58b0]]
= vserver object-store-server bucket lifecycle-management-rule show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the lifecycle management rule

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket lifecycle-management-rule show`
command displays information about lifecycle management rules. If no
parameters are specified, the command displays the following information
about all lifecycle management rule:

* Vserver name
* Bucket name
```

```
* Lifecycle Management rule identifier
* Lifecycle Management rule index
* Action
* Enabled



To display detailed information about a single lifecycle management rule,
run the command with the `-vserver``-bucket``-rule-id` and `-index`
parameters. The detailed view provides all of the information in the
default list and the following additional information:

* Link-id
* Prefix to be matched with object names
* Tags
* Minimum size of the object
* Maximum size of the object
* Lifecycle Management Rule Action
* Number of days since creation after which objects can be deleted
* Specific Date when the objects should expire cleanup object delete
markers
* Number of latest non-current versions to be retained
* Number of days after which non-current versions will be deleted
* Number of days of initiation after which upload can be aborted



To display detailed information about all lifecycle management rules, run
the command with the `-instance` parameter.

You can specify additional parameters to display information that matches
only those parameters. For example, to display information only about
lifecycle management rules with a specified rule identifier, run the
command with the `-rule-id` specified rule identifier parameter.

== Parameters

{ ``[-fields <fieldname>,...]``::
This specifies the fields that need to be displayed.

| ``[-instance  ]`` }::
If this parameter is specified, the command only displays information
about all lifecycle management rule entries.

``[-vserver <Vserver Name>]`` - Vserver Name::
If this parameter is specified, the command only displays information
about all lifecycle management rule entries on the specified vserver.
```

```
``[-bucket <TextNoCase>]`` - Object Store Server Bucket Name::
If this parameter is specified, the command only displays information
about all lifecycle management rule entries on the specified bucket.

``[-rule-id <text>]`` - Lifecycle Management Rule Identifier::
If this parameter is specified, the command only displays information
about all lifecycle management rule entries with the specified rule
identifier.

``[-index <integer>]`` - Lifecycle Management Rule Index::
If this parameter is specified, the command only displays information
about all lifecycle management rule entries with the specified index.

``[-is-enabled {true|false}]`` - Is This Rule Enabled?::
If this parameter is specified, the command only displays information
about lifecycle management rules that are enabled (true) or disabled
(false).

``[-prefix <text>]`` - Prefix to be Matched with Object Names::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified prefix.

``[-tags <text>,...]`` - Tags in Format <tag> or <tag=value>::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified tags.

``[-obj-size-greater-than {<integer>[KB|MB|GB|TB|PB]}]`` - Min Size of the
Object::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified object greater than
size.

``[-obj-size-less-than {<integer>[KB|MB|GB|TB|PB]}]`` - Max Size of the
Object::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified object lesser than
size.

``[-action
{Expiration|NoncurrentVersionExpiration|AbortIncompleteMultipartUpload|Exp
irationInternal}]`` - Lifecycle Management Action::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified action.

``[-obj-age-days <integer>]`` - Number of Days since Creation, After Which
```

```
Current Version of Objects Can be Deleted::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified number of days since
creation, after which objects can be deleted.

``[-obj-exp-date <MM/DD/YYYY HH:MM:SS>]`` - Specific Date When the Objects
Should Expire::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified date from when objects
can expire.

``[-expired-obj-del-marker {true|false}]`` - Cleanup Object Delete
Markers::
If this parameter is specified, the command only displays information
about lifecycle management rules that have delete markers enabled (true)
or disabled (false).

``[-new-non-curr-versions <integer>]`` - Number of Latest Non-current
Versions to Be Retained::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified non current versions
to be allowed.

``[-non-curr-days <integer>]`` - Number of Days after Which Non-current
Versions will Be Deleted::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified number of days after
which non-current versions can be deleted.

``[-after-initiation-days <integer>]`` - Number of Days of Initiation,
After Which Upload Can Be Aborted::
If this parameter is specified, the command only displays information
about lifecycle management rules with the specified number of days of
initiation, after which uploads can be aborted.

== Examples

The following example displays object store server bucket lifecycle
management rule entries for all Vservers.
```

cluster1::> vserver object-store-server bucket lifecycle-management-rule show
Vserver Bucket Rule-identifier Action-identifier Action Enabled
-------- --------- --------------- ---------------- -------------- ----------
vs1 bucket1 rule1 1 Expiration true
vs1 bucket1 rule2 1 AbortIncompleteMultipartUpload true
vs1 bucket2 rule1 1 Expiration true
vs1 bucket2 rule2 1 AbortIncompleteMultipartUpload true

vs1 bucket3 rule1 1 Expiration true
vs1 bucket3 rule2 1 AbortIncompleteMultipartUpload true
vs2 bucket4 rule1 1 Expiration true
vs2 bucket4 rule2 1 AbortIncompleteMultipartUpload true
vs2 bucket5 rule1 1 Expiration true
vs2 bucket5 rule2 1 AbortIncompleteMultipartUpload true
10 entries were displayed.

```
The following example displays detailed information about a lifecycle
management rule on bucket bucket1 on a Vserver named vs1 with rule
identifier rule1 and index 1:
```

cluster1::*> vserver object-store-server bucket lifecycle-management-rule show -vserver vs1 -bucket bucket1 -rule-id rule1 -index 1

Vserver: vs1
Object Store Server Bucket Name: bucket1
Lifecycle Management Rule Identifier: rule1
Lifecycle Management Rule Index: 1
Link-id from the Fabriclink Links Table: 4
Is This Rule Enabled?: true
Prefix to Be Matched with Object Names: obj1/
Tags in Format: -
Minimum Size of the Object: -
Maximum Size of the Object: -
Lifecycle Management Rule Action: Expiration
Number of Days Since Creation After Which
Objects Can Be Deleted: 100
Specific Date When the Objects Should Expire: -
Cleanup Object Delete Markers: -
Number of Latest Non-current Versions to be
Retained: -
Number of Days after Which Non-current
Versions will Be Deleted: -
Number of Days of Initiation After Which
Upload Can Be Aborted: -

```
[[ID4935c4a19807667d33f3900f62939f9b]]
= vserver object-store-server bucket policy-statement-condition create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a bucket policy statement condition

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.
```

```
== Description

The `vserver object-store-server bucket policy-statement-condition create`
command creates a single condition for a bucket policy statement in an
object store server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the Vserver on which the bucket
policy statement condition needs to be created for the object store server
bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the object store server bucket for
which the policy statement condition needs to be created. The object store
bucket must already exist.

``-index <integer>`` - Statement Index [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the index of the object store server bucket
policy statement in which a condition needs to be created. The index must
already exist.

``-operator {ip-address|not-ip-address|string-equals|string-not-
equals|string-equals-ignore-case|string-not-equals-ignore-case|string-
like|string-not-like|numeric-equals|numeric-not-equals|numeric-greater-
than|numeric-greater-than-equals|numeric-less-than|numeric-less-than-
equals}`` - Policy Condition Operator [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the condition operator to be applied on the
condition keys specified.

``[-source-ips <IP Address or Subnet>,...]`` - List of IP Addresses with
Access Allowed or Denied [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a list of IP addresses for which the access
is allowed or denied based on the operator specified.

``[-usernames <text>,...]`` - List of Usernames with Access Allowed or
Denied [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a list of object store server users for
which the access is allowed or denied based on the operator specified. The
user name policy variables '${aws:username}' can be specified here,
```

```
serving as placeholders that are dynamically replaced with the actual user
name during run time based on the request context.

``[-prefixes <text>,...]`` - List of Prefixes to be Matched [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify a list of prefixes that are compared with
the input prefix value specified at the time of execution of an S3-based
command, using the condition operator specified. The user name policy
variables '${aws:username}' can be specified here, serving as placeholders
that are dynamically replaced with the actual user name during run time
based on the request context.

``[-max-keys <integer>,...]`` - List of Maximum Keys Allowed to be Fetched
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a list of max-keys values that are allowed
or denied retrieval using an S3 list operation, based on the condition
operator specified.

``[-delimiters <text>,...]`` - List of Delimiters to be Matched [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify a list of delimiters that are compared with
the input delimiter value specified at the time of execution of an S3-
based command, using the condition operator specified. The user name
policy variables '${aws:username}' can be specified here, serving as
placeholders that are dynamically replaced with the actual user name
during run time based on the request context.

== Examples

The following example creates an object store server bucket policy
statement condition for storage virtual machine (SVM) vs1, bucket bucket1,
index 1 and ip-address as operator.
```

cluster1::*> vserver object-store-server bucket policy-statement-condition create -vserver vs1 -bucket bucket1 -index 1 -operator ip-address -source-ips 10.1.1.0/24,10.1.1.1

```
The following example creates an object store server bucket policy
statement condition for storage virtual machine (SVM) vs1, bucket bucket1,
index 1, string-like as operator and prefix with the user name policy
variable.
```

cluster1::*> vserver object-store-server bucket policy-statement-condition create -vserver vs1 -bucket bucket1 -index 1 -operator string-like -prefixes ${aws:username}/*

```
[[ID3b2fcdea4f3b1b4f5b9a3859525eec40]]
```

```
= vserver object-store-server bucket policy-statement-condition delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a bucket policy statement condition

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver object-store-server bucket policy-statement-condition delete`
command deletes a condition for the specified bucket policy statement
belonging to the object store server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the Vserver for which a condition
belonging to a particular bucket policy statement (which belongs to the
object store server bucket) you want to delete.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the object store server bucket for
which a condition belonging to a particular bucket policy statement needs
to be deleted.

``-index <integer>`` - Statement Index [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the index of the object store server bucket
policy for which a condition needs to be deleted.

``-operator {ip-address|not-ip-address|string-equals|string-not-
equals|string-equals-ignore-case|string-not-equals-ignore-case|string-
like|string-not-like|numeric-equals|numeric-not-equals|numeric-greater-
than|numeric-greater-than-equals|numeric-less-than|numeric-less-than-
equals}`` - Policy Condition Operator [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the condition operator of a condition which needs
to be deleted.

== Examples
```

```
The following example deletes an object store server bucket policy
statement condition for Vserver vs1, bucket bucket1, index 1 and operator
as IpAddress.
```

cluster1::*> vserver object-store-server bucket policy-statement-condition delete -vserver vs1 -bucket bucket1 -index 1 -operator IpAddress

```
[[ID52e2933824c10b7de7c7724d47ca57bc]]
= vserver object-store-server bucket policy-statement-condition modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a bucket policy statement condition

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver object-store-server bucket policy-statement-condition modify`
command modifies a single condition for a bucket policy statement in an
object store server bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the Vserver on which the bucket
policy statement condition needs to be modified for the object store
server bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the object store server bucket for
which the policy statement condition needs to be modified.

``-index <integer>`` - Statement Index [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the index of the object store server bucket
policy statement in which a condition needs to be modified.

``-operator {ip-address|not-ip-address|string-equals|string-not-
```

```
equals|string-equals-ignore-case|string-not-equals-ignore-case|string-
like|string-not-like|numeric-equals|numeric-not-equals|numeric-greater-
than|numeric-greater-than-equals|numeric-less-than|numeric-less-than-
equals}`` - Policy Condition Operator [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the condition operator to be applied on the
condition keys specified.

``[-source-ips <IP Address or Subnet>,...]`` - List of IP Addresses with
Access Allowed or Denied [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a list of IP addresses for which the access
is allowed or denied based on the operator specified.

``[-usernames <text>,...]`` - List of Usernames with Access Allowed or
Denied [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a list of object store server users for
which the access is allowed or denied based on the operator specified. The
user name policy variables '${aws:username}' can be specified here,
serving as placeholders that are dynamically replaced with the actual user
name during run time based on the request context.

``[-prefixes <text>,...]`` - List of Prefixes to be Matched [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify a list of prefixes that are compared with
the input prefix value specified at the time of execution of an S3-based
command, using the condition operator specified. The user name policy
variables '${aws:username}' can be specified here, serving as placeholders
that are dynamically replaced with the actual user name during run time
based on the request context.

``[-max-keys <integer>,...]`` - List of Maximum Keys Allowed to be Fetched
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify a list of max-keys values that are allowed
or denied retrieval using an S3 list operation, based on the condition
operator specified.

``[-delimiters <text>,...]`` - List of Delimiters to be Matched [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify a list of delimiters that are compared with
the input delimiter value specified at the time of execution of an S3-
based command, using the condition operator specified. The user name
policy variables '${aws:username}' can be specified here, serving as
placeholders that are dynamically replaced with the actual user name
during run time based on the request context.

== Examples
```

```
The following example modifies an object store server bucket policy
statement condition for storage virtual machine (SVM) vs1, bucket bucket1,
index 1 and ip-address as operator.
```

cluster1::*> vserver object-store-server bucket policy-statement-condition modify -vserver vs1 -bucket bucket1 -index 1 -operator ip-address -source-ips 10.1.0.0/16,10.1.1.1

```
The following example modifies an object store server bucket policy
statement condition for storage virtual machine (SVM) vs1, bucket bucket1,
index 1, string-like as operator and prefix with the user name policy
variable.
```

cluster1::*> vserver object-store-server bucket policy-statement-condition modify -vserver vs1 -bucket bucket1 -index 1 -operator string-like -prefixes "${aws:username}/*"

```
[[ID0a52f2fa7cbc3c892af0eb545111b2a9]]
= vserver object-store-server bucket policy-statement-condition show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the bucket policy condition

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver object-store-server bucket policy-statement-condition show`
command displays information about object store server bucket policy
condition.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.
```

```
``[-vserver <Vserver Name>]`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions for the specified
Vserver.

``[-bucket <TextNoCase>]`` - Object Store Server Bucket Name [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions for the specified
bucket.

``[-index <integer>]`` - Statement Index [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions for the specified
bucket policy index.

``[-operator {ip-address|not-ip-address|string-equals|string-not-
equals|string-equals-ignore-case|string-not-equals-ignore-case|string-
like|string-not-like|numeric-equals|numeric-not-equals|numeric-greater-
than|numeric-greater-than-equals|numeric-less-than|numeric-less-than-
equals}]`` - Policy Condition Operator [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions that match the
specified condition operator.

``[-source-ips <IP Address or Subnet>,...]`` - List of IP Addresses with
Access Allowed or Denied [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions that match the
specified bucket policy condition source IP addresses.

``[-usernames <text>,...]`` - List of Usernames with Access Allowed or
Denied [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions that match the
specified usernames.

``[-prefixes <text>,...]`` - List of Prefixes to be Matched [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions that match the
specified prefixes.
```

```
``[-max-keys <integer>,...]`` - List of Maximum Keys Allowed to be Fetched
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions that match the
specified max-keys.

``[-delimiters <text>,...]`` - List of Delimiters to be Matched [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information on the
object store server bucket policy statement conditions that match the
specified delimiters.

== Examples

The following example displays information on object store server bucket
policy statement conditions for vserver vs1, bucket bb1 and index 1:
```

cluster1::*> vserver object-store-server bucket policy-statement-condition show -vserver vs1 -bucket bb1 -index 1
Vserver: vs1
Bucket: bb1

Index Operator Source-IPs Usernames Prefixes Max-Keys Delimiters
----- ---------- ---------- ---------- ---------- -------- ----------
1 ip-address 1.1.1.0/24 - - - -
1 string-like - user1 pref - delim1
2 entries were displayed.

```
[[ID3a96c3b9c97a820e5db5d746a971cc9c]]
= vserver object-store-server bucket policy statement create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a bucket policy statement

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket policy statement create` command
creates a bucket policy statement for the object store server bucket.

== Parameters
```

```
``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the bucket
policy statement needs to be created for the object store server bucket.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which the policy statement needs to be created. The object store bucket
must already exist.

``[-index <integer>]`` - Statement Index::
This parameter specifies the index of the object store server bucket
policy statement. The allowed range is 1-10. This is an optional
parameter.

``-effect {deny|allow}`` - Allow or Deny Access::
Use this parameter to specify whether access is allowed or denied when a
user requests the specific action.

``[-action <Action>,...]`` - Bucket Policy Action Allowed or Denied::
Use this parameter to specify resource operations. The set of resource
operations that the object store server supports are GetObject, PutObject,
DeleteObject, ListBucket, GetBucketAcl, GetObjectAcl,
ListBucketMultipartUploads, ListMultipartUploadParts, GetObjectTagging,
PutObjectTagging, DeleteObjectTagging, PutBucketPolicy, GetBucketPolicy,
DeleteBucketPolicy, GetBucketLocation, GetBucketVersioning,
PutBucketVersioning, and ListBucketVersions. Wildcards are accepted for
this parameter.

``[-principal <Objectstore Principal>,...]`` - List of Users to Be Allowed
or Denied Access::
Validate the user requesting access against the object store server users
or groups or NAS groups specified in this parameter. To gain access, the
user in the context should either match one of the users or belong to one
of the groups specified in this principle parameter. An object store
server group is specified by adding a prefix "group/" to the group name. A
NAS group is specified by adding a prefix "nasgroup/" to the group name.

``[-resource <text>,...]`` - Bucket or Objects to Be Allowed or Denied
Access::
Use this parameter to specify the bucket, folder, or object for which
allow or deny permissions are set. The user name policy variables
'${aws:username}' can be specified here, serving as placeholders that are
dynamically replaced with the actual user name during run time based on
the request context.
```

```
``[-sid <SID>]`` - Statement Identifier::
This optional parameter specifies a text comment for the object store
server bucket policy statement. Alpha numeric characters are allowed as
values for this parameter.

== Examples

The following example creates an object store server bucket policy
statement for storage virtual machine (SVM) vs1 and bucket1 which
specifies allowed access to a readme folder for the object store server
user user1.
```

cluster1::> vserver object-store-server bucket policy statement create -vserver vs1 -bucket bucket1 -effect allow -action GetObject,PutObject,DeleteObject,ListBucket -principal user1,group/group1,nasgroup/group2 -resource bucket1/readme/* -sid "fullAccessToReadmeForUser1"

```
The following example creates an object store server bucket policy
statement for storage virtual machine (SVM) vs1 and bucket1 which
specifies allowed access to the corresponding user home directory by
specifying the user name policy variable in the resource field.
```

cluster1::> vserver object-store-server bucket policy statement create -vserver vs1 -bucket bucket1 -effect allow -action * -principal * -resource bucket1,bucket1/${aws:username}/* -sid "fullAccessToUsersHomeDirectory"

```
[[IDafd590d8576c5e158adb1c77c515e8da]]
= vserver object-store-server bucket policy statement delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a bucket policy statement

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket policy statement delete` command
deletes the bucket policy statement belonging to the object store server
bucket.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver whose bucket policy
statement (which belongs to the object store server bucket) you want to
delete.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket whose
policy needs to be deleted.

``-index <integer>`` - Statement Index::
This parameter specifies the index of the object store server bucket
policy.

== Examples

The following example deletes an object store server bucket policy
statement with index 1 of Vserver vs1 and bucket bucket1.
```

cluster1::> vserver object-store-server bucket policy statement delete -vserver vs1 -bucket bucket1 -index 1

```
[[IDe73af5fa6bc6dbe56c4edd87791b8884]]
= vserver object-store-server bucket policy statement modify
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a bucket policy statement

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket policy statement modify` command
modifies a bucket policy statement.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server bucket for which the bucket policy statement needs to be modified.

``-bucket <TextNoCase>`` - Object Store Server Bucket Name::
This parameter specifies the name of the object store server bucket for
which policy statement needs to be modified.

``-index <integer>`` - Statement Index::
This parameter specifies the index of the object store server bucket
policy statement.

``[-effect {deny|allow}]`` - Allow or Deny Access::
Use this parameter to specify whether access is allowed or denied when a
user requests the specific action.

``[-action <Action>,...]`` - Bucket Policy Action Allowed or Denied::
Use this parameter to specify resource operations. The set of resource
operations that the object store server supports are GetObject, PutObject,
DeleteObject, ListBucket, GetBucketAcl, GetObjectAcl,
ListBucketMultipartUploads, ListMultipartUploadParts, GetObjectTagging,
PutObjectTagging, DeleteObjectTagging, GetBucketLocation, PutBucketPolicy,
GetBucketPolicy, DeleteBucketPolicy, GetBucketVersioning,
PutBucketVersioning, and ListBucketVersions.

``[-principal <Objectstore Principal>,...]`` - List of Users to Be Allowed
or Denied Access::
Validate the user requesting access against the object store server users
or groups or NAS groups specified in this parameter. To gain access, the
user in the context should either match one of the users or belong to one
of the groups specified in this principle parameter. An object store
```

```
server group is specified by adding a prefix "group/" to the group name. A
NAS group is specified by adding a prefix "nasgroup/" to the group name.

``[-resource <text>,...]`` - Bucket or Objects to Be Allowed or Denied
Access::
Use this parameter to specify the bucket, folder, or object for which
allow or deny permissions are set. The user name policy variables
'${aws:username}' can be specified here, serving as placeholders that are
dynamically replaced with the actual user name during run time based on
the request context.

``[-sid <SID>]`` - Statement Identifier::
This optional parameter specifies a text comment for the object store
server bucket policy statement.

== Examples

The following example modifies an object store server bucket policy
statement for storage virtual machine (SVM) vs1 and bucket1 which
specifies allowed access to a readme folder for the object store server
user user1.
```

cluster1::> vserver object-store-server bucket policy statement modify -vserver vs1 -bucket bucket1 -index 1 -effect allow -action GetObject,PutObject,DeleteObject,ListBucket -principal user1 -resource bucket1/readme/* -sid "fullAccessToReadmeForUser1"

```
The following example modifies an object store server bucket policy
statement for storage virtual machine (SVM) and bucket1 which specifies
allowed access to the corresponding user home directory by specifying the
user name policy variable in the resource field.
```

cluster1::> vserver object-store-server bucket policy statement modify -vserver vs1 -bucket bucket1 -index 1 -effect allow -action * -principal * -resource bucket1,bucket1/${aws:username}/* -sid "fullAccessToUsersHomeDirectory"

```
[[ID0ddbefeea1fc126abfa465b936808f33]]
= vserver object-store-server bucket policy statement show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show the bucket policy
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket policy statement show` command
displays information about object store server bucket policy.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information on the
object store server bucket policy statements for the specified Vserver.

``[-bucket <TextNoCase>]`` - Object Store Server Bucket Name::
If you specify this parameter, the command displays information on the
object store server bucket policy statements for the specified bucket.

``[-index <integer>]`` - Statement Index::
If you specify this parameter, the command displays information on the
object store server bucket policy statements that match the specified
index.

``[-effect {deny|allow}]`` - Allow or Deny Access::
If you specify this parameter, the command displays information on the
object store server bucket policy statements that match the specified
effect.

``[-action <Action>,...]`` - Bucket Policy Action Allowed or Denied::
If you specify this parameter, the command displays information on the
object store server bucket policy statements that match the specified
action.

``[-principal <Objectstore Principal>,...]`` - List of Users to Be Allowed
or Denied Access::
If you specify this parameter, the command displays information on the
object store server bucket policy statements that match the specified
bucket principal.
```

```
``[-resource <text>,...]`` - Bucket or Objects to Be Allowed or Denied
Access::
If you specify this parameter, the command displays information on the
object store server bucket policy statements that match the specified
resource.

``[-sid <SID>]`` - Statement Identifier::
If you specify this parameter, the command displays information on the
object store server bucket policy statements that match the specified sid.

== Examples

The following example displays information on object store server bucket
policy statements for vserver vs1 and bucket bucket1:
```

cluster1::> vserver object-store-server bucket policy show -vserver vs1 -bucket bucket1
Vserver Bucket Index Effect Action Principal Resource
---------- ---------- ----- ------ ---------- --------- ---------
vs1
bucket1 1 allow GetObject, user1 bucket1/
PutObject, readme/*
DeleteObject
, ListBucket
bucket1 2 allow GetObject user2 bucket1/*
2 entries were displayed.

```
The following example displays detailed information of the object store
server bucket policy statement associated with Vserver vs1 and bucket
bucket1:
```

cluster1::> vserver object-store-server bucket policy show -vserver vs1 -bucket bucket1 -index 1
Vserver :vs1
Bucket :bucket1
Index :1
Effect :allow
Action :GetObject
Principal :user-2
Resource :bucket1/readme/*
Sid :AllowAccessToUseruser1ForGetObject

```
[[IDdc26de906e22ee3f365feeb0fb891181]]
= vserver object-store-server bucket snapshot create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Create a snapshot for the S3 bucket.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket snapshot create` command creates a
snapshot of a specified S3 bucket.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver that contains the S3 bucket for
which the snapshot is to be created.

``-bucket <TextNoCase>`` - S3 Bucket Name::
This specifies the name of the S3 bucket where the snapshot is to be
created.

``-snapshot <TextNoCase>`` - Snapshot Name::
This specifies the name of the snapshot that is to be created.

``[-foreground <true>]`` - Run Operation in the Foreground::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). If true, the
command will wait until the operation is finished. If false, the command
will return immediately and the operation will complete in the background.

== Examples





[[ID245e895fc9e339cea0c9415a5b60226a]]
= vserver object-store-server bucket snapshot delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a snapshot for the S3 bucket.
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server bucket snapshot delete` command deletes a
snapshot from a specified S3 bucket.

== Parameters

``-vserver <vserver name>`` - Vserver::
This specifies the name of the Vserver that contains the S3 bucket for
from which the snapshot will be deleted.

``-bucket <TextNoCase>`` - S3 Bucket Name::
This specifies the name of the S3 bucket for from which the snapshot will
be deleted.

``-snapshot <TextNoCase>`` - Snapshot Name::
This specifies the name of the snapshot that is to be deleted.

``[-foreground <true>]`` - Run Operation in the Foreground::
This specifies whether the operation runs in the foreground. The default
setting is `true` (the operation runs in the foreground). If true, the
command will wait until the operation is finished. If false, the command
will return immediately and the operation will complete in the background.

== Examples





[[ID5e4cf6dbdf00ec5120b11e89885b311d]]
= vserver object-store-server bucket snapshot show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of snapshots associated with the S3 bucket.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver object-store-server bucket snapshot show` command displays
information about snapshots created for S3 buckets.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the snapshots located for the specified Vserver.

``[-bucket <TextNoCase>]`` - S3 Bucket Name::
If you specify this parameter, the command displays information only about
the snapshots located for the specified S3 buckets.

``[-snapshot <TextNoCase>]`` - Snapshot Name::
If you specify this parameter, the command displays information only
snapshots matching the specified name.

``[-create-time <Date>]`` - Snapshot Creation Time::
If this parameter is specified, the command displays information only for
the snapshots that match the specified creation time.

``[-instance-uuid <UUID>]`` - Snapshot Instance UUID::
If this parameter is specified, the command displays information only for
the snapshots that have the specified instance UUID.

== Examples





[[IDe74008cc8099b9409c68d1a8877d32d6]]
= vserver object-store-server group create
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an Object Store Server Group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server group create` command creates an object
store group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which to create the
object store group. The Vserver must already exist.

``-gid <integer>`` - Group ID::
This parameter specifies a unique ID used to identify a particular object
store group.

``-name <TextNoCase>`` - Group Name::
This parameter specifies the name of the object store group.

``[-users <TextNoCase>,...]`` - List of Users Belonging to the Group::
Use this parameter to specify the list of object store users who belong to
the object store group.

``[-policies <TextNoCase>,...]`` - List of Policies Attached to the
Group::
Use this parameter to specify the list of object store policies that are
attached to the object store group.

``[-comment <text>]`` - Group Description::
This optional parameter specifies a text comment for the object store
group.

== Examples

The following example creates an object store group named user_group for
Vserver vs1:
```

cluster1::> vserver object-store-server group create -vserver vs1 -name user_group -users user1,user2 -policies policy1,policy2 -comment "UserGroup1"

```
[[IDcc6376352c456e17fbf910f5e47539bb]]
= vserver object-store-server group delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an Object Store Server Group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server group delete` command deletes an object
store group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server you want to delete.

``-gid <integer>`` - Group ID::
This parameter specifies the ID of the object store group you want to
delete.

== Examples

The following example deletes an object store group for Vserver vs1:
```

cluster1::> vserver object-store-server group delete -vserver vs1 -gid 1

```
[[ID4a84aa384a9986bd9c4512772640e417]]
= vserver object-store-server group modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an Object Store Server Group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server group modify` command modifies an object
store group.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
group which you want to modify.

``-gid <integer>`` - Group ID::
This parameter specifies the ID of the object store group.

``[-name <TextNoCase>]`` - Group Name::
This parameter specifies the name of the object store group.

``[-users <TextNoCase>,...]`` - List of Users Belonging to the Group::
Use this parameter to specify the list of object store users who belong to
the object store group.

``[-policies <TextNoCase>,...]`` - List of Policies Attached to the
Group::
Use this parameter to specify the list of object store policies that are
attached to the object store group.

``[-comment <text>]`` - Group Description::
This parameter specifies the text comment for the object store group.

== Examples

The following example modifies the comment of the object store group for
Vserver vs1:
```

cluster1::> vserver object-store-server group modify -vserver vs1 -gid 3 -comment "UserGroup"

```
[[IDd87796e533828ac3d12fc904b8075c37]]
= vserver object-store-server group show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Object Store Server Groups

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server group show` command displays information
about the object store group.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields <fieldname>, ...` parameter, the command
output also includes the specified field or fields. You can use '-fields
?' to display the fields to specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information on the
object store server groups for the specified Vserver.

``[-gid <integer>]`` - Group ID::
If you specify this parameter, the command displays information on the
object store server group that match the specified group ID.

``[-name <TextNoCase>]`` - Group Name::
If you specify this parameter, the command displays information on the
object store server groups that match the specified group name.

``[-users <TextNoCase>,...]`` - List of Users Belonging to the Group::
If you specify this parameter, the command displays information on the
object store server groups that match the specified user.
```

```
``[-policies <TextNoCase>,...]`` - List of Policies Attached to the
Group::
If you specify this parameter, the command displays information on the
object store server groups that match the specified policy.

``[-comment <text>]`` - Group Description::
If you specify this parameter, the command displays information on the
object store server groups that match the specified comment.

== Examples

The following example displays information for all object store groups in
admin privilege:
```

cluster1::> vserver object-store-server group show
Vserver Group ID Group Name Users Policies
---------- --------- ---------- ---------- ----------
vs1 3 UserGroup user1, user2 policy1, policy2
Comment: User_Privileges
vs1 4 AdminGroup admin1, admin2 policy1, policy2
Comment: Admin_Privileges
2 entries were displayed.

```
The following example displays information for a particular object store
group associated with vserver vs1:
```

cluster1::> vserver object-store-server group show -vserver vs1 -gid 5
Vserver Name :vs1
Group ID :5
Group Name :User-Group
Users :user_1, user_2
Policies :Policy1, Policy2, Policy3
Comment :User group

```
[[IDbffc115b0f736dc50e900a8135c60d2b]]
= vserver object-store-server policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server policy create` command creates an object
store policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which to create the
object store policy. The Vserver must already exist.

``-policy <TextNoCase>`` - Policy Name::
This parameter specifies the name of the object store policy.

``[-comment <text>]`` - Comment::
This optional parameter specifies a text comment for the object store
policy.

== Examples

The following example creates an object store policy named Policy_1 for
Vserver vs1:
```

cluster1::> vserver object-store-server policy create -vserver vs1 -policy Policy_1 -comment "ReadAccessForBucket1"

```
[[IDa203b0824d16aa9870314c418f47d083]]
= vserver object-store-server policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server policy delete` command deletes an object
store policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server you want to delete.

``-policy <TextNoCase>`` - Policy Name::
This parameter specifies the name of the object store policy you want to
delete.

== Examples

The following example deletes an object store policy for Vserver vs1:
```

cluster1::>vserver object-store-server policy delete -vserver vs1 -policy Policy_2

```
[[IDacc4a92d9f31b17f087eb89a2da53b49]]
= vserver object-store-server policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server policy modify` command modifies an object
store policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
policy you want to modify.

``-policy <TextNoCase>`` - Policy Name::
This parameter specifies the name of the object store policy.

``[-comment <text>]`` - Comment::
This parameter specifies the text comment for the object store policy.

== Examples

The following example modifies the comment of the object store policy for
Vserver vs1:
```

cluster1::> vserver object-store-server policy modify -policy Policy_1 -comment "Read_Access_for_Bucket2"

```
[[IDf2131d0fdbb9f8c32d5948941a5fe567]]
= vserver object-store-server policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Show the policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server policy show` command displays information
about the object store policy.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information on the
object store server policies for the specified Vserver.

``[-policy <TextNoCase>]`` - Policy Name::
If you specify this parameter, the command displays information on the
object store server policies that match the specified policy name.

``[-is-read-only {true|false}]`` - Is Read-Only?::
If you specify this parameter, the command displays information on the
object store server policies that match the specified read only field.

``[-comment <text>]`` - Comment::
If you specify this parameter, the command displays information on the
object store server policies that match the specified comment.

== Examples

The following example displays information for all object store policies
in admin privilege:
```

```
cluster1::> vserver object-store-server policy show
Vserver Name Is Read-Only Comment
----------- ----------------- ------------ ----------------
vs1 FullAccess true Read Only Policy: To allow full access to S3 resources
vs1 NoS3Access true Read Only Policy: To deny access to all S3 resources
```

vs1 Policy_1 false Read_access_for_bucket1
vs1 Policy_2 false Read_access_for_bucket2
vs1 ReadOnlyAccess true Read Only Policy: To allow read-only access to S3 resources
5 entries were displayed.

```
The following example displays information for a particular object store
policy associated with Vserver vs1:
```

cluster1::> vserver object-store-server policy show -policy Policy_1
Vserver Name Is Read-Only Comment
---------- ---------- ---------- ----------
vs1 Policy_1 false Read_access_for_bucket1

```
[[ID58e74f50a082307f5e9d380416194848]]
= vserver object-store-server policy statement create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a Policy Statement

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server policy statement create` command creates
a policy statement for the object store server policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the policy
statement needs to be created for the object store server policy.

``-policy <TextNoCase>`` - Policy Name::
This parameter specifies the name of the object store server policy for
which the policy statement needs to be created. The object store policy
must already exist.

``-index <integer>`` - Statement Index::
This parameter specifies the unique index used to identify the particular
object store server policy statement.
```

```
``-effect {deny|allow}`` - Allow or Deny Access::
Use this parameter to specify whether or not access is allowed or denied
when a user requests a specific action.

``-actions <Action>,...`` - Policy Actions::
Use this parameter to specify resource operations. The set of resource
operations that the object store server supports are GetObject, PutObject,
DeleteObject, ListBucket, GetBucketAcl, GetObjectAcl, ListAllMyBuckets,
ListBucketMultipartUploads, ListMultipartUploadParts, GetObjectTagging,
PutObjectTagging, DeleteObjectTagging, CreateBucket, DeleteBucket,
GetBucketLocation, GetBucketVersioning, PutBucketVersioning, and
ListBucketVersions. Wildcards are accepted for this parameter. If all
operations must be specified, then use the wildcard character ``_*_`` to
specify it. The default actions are GetObject, PutObject, DeleteObject,
ListBucket, GetBucketAcl, GetObjectAcl, ListBucketMultipartUploads,
ListMultipartUploadParts, GetObjectTagging, PutObjectTagging,
DeleteObjectTagging, GetBucketLocation, PutBucketPolicy, GetBucketPolicy,
DeleteBucketPolicy, GetBucketVersioning, PutBucketVersioning, and
ListBucketVersions.

``-resource <text>,...`` - Buckets or Objects::
Use this parameter to specify the bucket, folder, or object for which
allow or deny permissions are set. The user name policy variables
'${aws:username}' can be specified here, serving as placeholders that are
dynamically replaced with the actual user name during run time based on
the request context.

``[-sid <SID>]`` - Statement Identifier::
This optional parameter specifies a text comment for the object store
server policy statement. The parameter name "sid" refers to statement
identifier.

== Examples

The following example creates an object store server policy statement for
storage virtual machine (SVM) vs1 and Policy_1 which specifies allowed
access to bucket1 resources.
```

cluster1::> vserver object-store-server policy statement create -vserver vs1 -policy Policy_1 -effect allow -actions GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetObjectAcl,ListAllMyBuckets,GetBucketLocation -resource bucket1/* -sid "FullAccesToBucket1"

```
The following example creates an object store server policy statement for
storage virtual machine (SVM) vs1 and Policy_1 which specifies allowed
access to the corresponding user home directory by specifying the user
name policy variable in the resource field.
```

cluster1::> vserver object-store-server policy statement create -vserver vs1 -policy Policy_1 -effect allow -actions * -resource bucket1,bucket1/${aws:username}/* -sid "fullAccessToUsersHomeDirectory"

```
[[IDfbc7a083c16afb723e8a5a7b46cf8435]]
= vserver object-store-server policy statement delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Policy Statement

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server policy statement delete` command deletes
the policy statement belonging to the object store server policy.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver whose policy statement
you want to delete.

``-policy <TextNoCase>`` - Policy Name::
This parameter specifies the name of the object store server policy whose
policy statement needs to be deleted.

``-index <integer>`` - Statement Index::
This parameter specifies the index of the object store server policy
statement.

== Examples

The following example deletes an object store server policy statement with
index 1 of Vserver vs1 and policy Policy_1.
```

cluster1::> vserver object-store-server policy statement delete -vserver vs1 -policy Policy_1 -index 1

```
[[ID5c1bbc5902dfeb2c2154868bf380e4a3]]
= vserver object-store-server policy statement modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a Policy Statement

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server policy statement modify` command modifies
a policy statement.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server policy for which the policy statement needs to be modified.

``-policy <TextNoCase>`` - Policy Name::
This parameter specifies the name of the object store server policy for
the policy statement that needs to be modified.

``-index <integer>`` - Statement Index::
This parameter specifies the index of the object store server policy
statement.

``[-effect {deny|allow}]`` - Allow or Deny Access::
Use this parameter to specify whether or not access is allowed or denied
when a user requests a specific action.

``[-actions <Action>,...]`` - Policy Actions::
Use this parameter to specify resource operations. The set of resource
operations that the object store server supports are GetObject, PutObject,
DeleteObject, ListBucket, GetBucketAcl, GetObjectAcl, ListAllMyBuckets,
ListBucketMultipartUploads, ListMultipartUploadParts, GetObjectTagging,
PutObjectTagging, DeleteObjectTagging, CreateBucket, DeleteBucket,
GetBucketLocation, PutBucketPolicy, GetBucketPolicy, DeleteBucketPolicy,
GetBucketVersioning and PutBucketVersioning.
```

```
``[-resource <text>,...]`` - Buckets or Objects::
Use this parameter to specify the bucket, folder, or object for which
allow or deny permissions are set. The user name policy variables
'${aws:username}' can be specified here, serving as placeholders that are
dynamically replaced with the actual user name during run time based on
the request context.

``[-sid <SID>]`` - Statement Identifier::
This optional parameter specifies a text comment for the object store
server policy statement.

== Examples

The following example modifies an object store server policy statement for
storage virtual machine (SVM) vs1 and Policy_1 which specifies allowed
access to bucket1 resources.
```

cluster1::> vserver object-store-server policy statement modify -vserver vs1 -policy Policy_1 -index 5 -effect allow -actions GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetObjectAcl,ListAllMyBuckets,GetBucketLocation -resource bucket1/* -sid FullAccesToBucket1Resources

```
The following example modifies an object store server policy statement for
storage virtual machine (SVM) vs1 and Policy_1 which specifies allowed
access to the corresponding user home directory by specifying the user
name policy variable in the resource field.
```

cluster1::> vserver object-store-server policy statement modify -vserver vs1 -policy Policy_1 -index 5 -effect allow -actions * -resource bucket1,bucket1/${aws:username}/* -sid "fullAccessToUsersHomeDirectory"

```
[[ID62e77642cecebcbb28af433b5cee28aa]]
= vserver object-store-server policy statement show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Show Policy Statements

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `vserver object-store-server policy statement show` command displays
information about object store server policy statements.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information on the
object store server policy statements for the specified Vserver.

``[-policy <TextNoCase>]`` - Policy Name::
If you specify this parameter, the command displays information on the
object store server policy statements for the specified policy.

``[-index <integer>]`` - Statement Index::
If you specify this parameter, the command displays information on the
object store server policy statements that match the specified index.

``[-effect {deny|allow}]`` - Allow or Deny Access::
If you specify this parameter, the command displays information on the
object store server policy statements that match the specified effect.

``[-actions <Action>,...]`` - Policy Actions::
If you specify this parameter, the command displays information on the
object store server policy statements that match the specified action.

``[-resource <text>,...]`` - Buckets or Objects::
If you specify this parameter, the command displays information on the
object store server policy statements that match the specified resource.

``[-sid <SID>]`` - Statement Identifier::
If you specify this parameter, the command displays information on the
object store server policy statements that match the specified sid.

== Examples

The following example displays information on object store server policy
statements for Vserver vs1 and policy Policy_1:
```

```
cluster1::> vserver object-store-server policy statement show -vserver vs1 -policy Policy_1
Vserver Policy Index Effect Actions Resources
--------- ------------ ------ ------ ---------------- -----------------
vs1
Policy_1 1 allow ListBucket *
Policy_1 5 allow GetObject, bucket1/*
PutObject,
DeleteObject,
ListBucket,
GetBucketAcl,
GetObjectAcl,
ListAllMyBuckets
Sid: FullAccesToBucket1
2 entries were displayed.
```

```
The following example displays detailed information of the object store
server policy statement associated with Vserver vs1 and policy Policy_1:
```

```
cluster1::> vserver object-store-server policy statement show -vserver vs1 -policy Policy_1 -index 5
Vserver: vs1
Policy: Policy_1
Index: 5
Effect: allow
Actions: GetObject, PutObject, DeleteObject, ListBucket, GetBucketAcl, GetObjectAcl, ListAllMyBuckets
Resource: bucket1/*
Sid: FullAccesToBucket1
```

```
[[IDa3ef76ee70c324b93e60c169d4b915c0]]
= vserver object-store-server user create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an object store server user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server user create` command creates an object
store user. This will generate an access-key and a secret-key to be used
for aws v4 authentication. The user keys generated will be shown as part
of the response of this command. Note that the secret-key is not
retrievable and should be noted down. It will not be shown as part of the
xref:{relative_path}vserver-object-store-server-user-show.html[vserver
```

```
object-store-server user show] command. If the secret-key is lost, the
user keys will need to be regenerated.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which to create the
object store user. The Vserver must already exist.

``-user <TextNoCase>`` - Object Store Server User Name::
This parameter specifies the name of the object store user. If user is a
part of Active Directory, user must be specified in User Principal Name
(UPN) format.

``[-comment <text>]`` - Object Store Server User Description::
This optional parameter specifies a text comment for the object store
user.

``[-access-key <text>]`` - Access Key for the Object Store Server User::
This optional parameter specifies the access key for the object store
user.

``[-secret-key <text>]`` - Secret Key for the Object Store Server User::
This optional parameter specifies the secret key for the object store
user.

``[-key-time-to-live {P[<integer>D]T[<integer>H][<integer>M][<integer>S] |
P<integer>W}]`` - Time Period After Which User Keys Expire::
This optional parameter specifies a time period after which the object
store user keys expire and are no longer valid. If the value specified is
zero or no value is specified, then the user keys will not expire.

== Examples

The following example creates an object store user user1 for Vserver vs1.
```

cluster1::> vserver object-store-server user create -vserver vs1 -user user1
Vserver: vs1
User: user1
Access Key: DUMMY_ACCESS_KEY_123
Secret Key: dummy_secret_key_1234_abcd__lkfj
Warning: The secret key won't be displayed again. Save this key for future use.

```
The following example creates an object store user user1, with the
specified keys, for Vserver vs1.
```

cluster1::> vserver object-store-server user create -vserver vs1 -user user1 -access-key DUMMY_ACCESS_KEY_123 -secret-key dummy_secret_key_1234_abcd*lkfj*
Vserver: vs1
User: user1
Access Key: DUMMY_ACCESS_KEY_123
*Secret Key: dummy_secret_key_1234_abcd*lkfj
Warning: The secret key won't be displayed again. Save this key for future use.

```
The following example creates an object store user user1, for Vserver vs1,
where user1 is a member of active directory.
```

cluster1::> vserver object-store-server user create -vserver vs1 -user user1@domain1.com -access-key DUMMY_ACCESS_KEY_123 -secret-key dummy_secret_key_1234_abcd*lkfj*
Vserver: vs1
User: user1@domain1.com
Access Key: DUMMY_ACCESS_KEY_123
*Secret Key: dummy_secret_key_1234_abcd*lkfj
Warning: The secret key won't be displayed again. Save this key for future use.

```
== Related Links

* xref:{relative_path}vserver-object-store-server-user-show.html[vserver
object-store-server user show]




[[IDa0be04bed42e84817c2e4de32709e0dd]]
= vserver object-store-server user delete-keys
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete keys for an object store user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server user delete-keys` command deletes the
access-key and secret-key for an object store user. To regenerate the keys
for the user again, use the xref:{relative_path}vserver-object-store-
server-user-regenerate-keys.html[vserver object-store-server user
regenerate-keys] command.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the SVM on which the keys should be
deleted for the object store user. The object store user must already
exist.

``-user <TextNoCase>`` - Object Store Server User Name::
This parameter specifies the name of the object store user.

== Examples

The following example deletes the keys for object store user for the SVM
vs1.
```

cluster1::> vserver object-store-server user delete-keys -vserver vs1 -user user1

```
== Related Links

* xref:{relative_path}vserver-object-store-server-user-regenerate-
keys.html[vserver object-store-server user regenerate-keys]




[[ID3f8fb96050f57f95ef0420e5ece66437]]
= vserver object-store-server user delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an object store server user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server user delete` command deletes an object
store user.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver for the object store
server you want to delete.

``-user <TextNoCase>`` - Object Store Server User Name::
This parameter specifies the name of the object store user you want to
delete.

== Examples

The following example deletes an object store user for Vserver vs1.
```

cluster1::> vserver object-store-server user delete -vserver vs1 -user user1

```
[[ID833b57c8d447e71002602e9eacb4d434]]
= vserver object-store-server user modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an object store server user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server user modify` command modifies an object
store user.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the SVM for the object store user
which you want to modify.

``-user <TextNoCase>`` - Object Store Server User Name::
This parameter specifies the name of the object store user.

``[-comment <text>]`` - Object Store Server User Description::
This parameter specifies the text comment for the object store user.

== Examples

The following example modifies the object store user for the SVM vs1:
```

cluster1::> vserver object-store-server user modify -vserver vs1 -user user1 -comment testuser --key-time-to -live 1h20m

```
[[ID9a4b14ec063c4e175bb52b6a16d66a1b]]
= vserver object-store-server user regenerate-keys
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Regenerate keys for object store user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server user regenerate-keys` command regenerates
a new access-key and secret-key for an object store user. The user keys
generated will be shown as part of the command response. Note that the
secret-key is not retrievable and should be noted down. It will not be
shown as part of the xref:{relative_path}vserver-object-store-server-user-
show.html[vserver object-store-server user show] command. If the secret-
key is lost, the user keys will need to be regenerated.

== Parameters

``-vserver <Vserver Name>`` - Vserver Name::
This parameter specifies the name of the Vserver on which the keys should
be generated for the object store user. The object store user must already
exist.

``-user <TextNoCase>`` - Object Store Server User Name::
This parameter specifies the name of the object store user.

``[-key-time-to-live {P[<integer>D]T[<integer>H][<integer>M][<integer>S] |
P<integer>W}]`` - Time Period after Which User Keys Expire::
This optional parameter specifies a time period after which the object
store user keys expire and are no longer valid. If the value specified is
zero, then the user keys will not expire.

``[-access-key <text>]`` - Access Key for the Object Store User::
This optional parameter specifies the access key for the object store
user.

``[-secret-key <text>]`` - Secret Key for the Object Store User::
This optional parameter specifies the secret key for the object store
user.

== Examples

The following example regenerates the keys for object store user for
Vserver vs1.
```

cluster1::> vserver object-store-server user regenerate-keys -vserver vs1 -user user1
Vserver: vs1

User: user1
Access Key: DUMMY_ACCESS_KEY_123
Secret Key: dummy_secret_key_1234_abcd__lkfj
Warning: The secret key won’t be displayed again. Save this key for future use.

```
The following example sets the new keys, that are passed from the input,
for object store user for Vserver vs1.
```

cluster1::> vserver object-store-server user regenerate-keys -vserver vs1 -user user1 -access-key DUMMY_ACCESS_KEY_123 -secret-key dummy_secret_key_1234_abcd*lkfj*
Vserver: vs1
User: user1
Access Key: DUMMY_ACCESS_KEY_123
*Secret Key: dummy_secret_key_1234_abcd*lkfj
Warning: The secret key won’t be displayed again. Save this key for future use.

```
== Related Links

* xref:{relative_path}vserver-object-store-server-user-show.html[vserver
object-store-server user show]




[[ID350468313b5173b35a17f162d8049cf5]]
= vserver object-store-server user show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display object store server users

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver object-store-server user show` command displays information
about the object store user. The user secret-key will not be shown as part
of the response. It is shown only when user is created or user keys are
regenerated.

== 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.

``[-vserver <Vserver Name>]`` - Vserver Name::
If you specify this parameter, the command displays information only about
the object store users for the specified SVM.

``[-user <TextNoCase>]`` - Object Store Server User Name::
If you specify this parameter, the command displays information only for
object store users that match the specified object store user name.

``[-id <integer>]`` - Object Store Server User ID::
If you specify this parameter, the command displays information only for
object store users that match the specified user id.

``[-comment <text>]`` - Object Store Server User Description::
If you specify this parameter, the command displays information only for
object store users that match the specified comment field.

``[-access-key <text>]`` - Access Key for the Object Store Server User::
If you specify this parameter, the command displays information only for
object store users that match the specified access key.

``[-key-time-to-live {P[<integer>D]T[<integer>H][<integer>M][<integer>S] |
P<integer>W}]`` - Time Period After Which User Keys Expire::
If you specify this parameter, the command displays information only for
object store users that match the specified key time to live value.

``[-key-expiry-time {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm]}]`` - Date And Time
When User Keys Expire::
If you specify this parameter, the command displays information only for
object store users that match the specified key expiry date value.

== Examples

The following example displays information of all object store users in
admin privilege:
```

```
cluster1::> vserver object-store-server user show
Vserver User ID Key Time To Live Key Expiry Time
---------- ------------------ -------- --------------- -----------------
```

vs1 user1 1 1h 2/6/2023 13:28:50
Comment: testuser
vs1 user2 2 - -
vs1 user@domain1.com 3 - -
2 entries were displayed.

```
The following example displays information of a particular object store
user associated with the SVM vs1 in admin privilege:
```

cluster1::> vserver object-store-server user show -vserver vs1 -user user1
Vserver Name: vs1
Object Store Server User Name: user1
Object Store Server User ID: 1
Object Store Server User Description: testuser
Access Key for the Object Store Server User: 5HBRV20PWWX7IIHKYRRN
Time Period After Which User Keys Expire: 1h
Date And Time When User Keys Expire: 2/6/2023 13:28:50

```
The following example displays information of all object store users in
advanced privilege:
```

cluster1::*> vserver object-store-server user show
Vserver User ID Key Time To Live Key Expiry Time
--- --- --- --- ---
vs1 root 0 - -
Access Key: -
Comment: Root User
vs1 user2 2 99999h0m0s 7/6/2034 23:40:54
Access Key: 2K4PL22JQV5L2WA564TB
Comment:
2 entries were displayed.

```
The following example displays information of a particular object store
user associated with the SVM vs1 in advanced privilege. Note that user
secret key is not shown as part of this command:
```

cluster1::*> vserver object-store-server user show -vserver vs1 -user user1
Vserver Name: vs1
Object Store Server User Name: user1
Object Store Server User ID: 1
Object Store Server User Description: testuser
Access Key for the Object Store Server User: 5HBRV20PWWX7IIHKYRRN
Time Period After Which User Keys Expire: 1h
Date And Time When User Keys Expire: 2/6/2023 13:28:50

```
:leveloffset: -1
```

```
= vserver peer commands

:leveloffset: +1


[[ID098be6eff0517f8e44798b717e47ef71]]
= vserver peer accept
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Accept a pending Vserver peer relationship

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer accept` command is used to accept the Vserver peer
relationship between two Vservers. This command is used only for
intercluster Vserver peer relationships.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Specifies name of the local Vserver for which you want to accept the
Vserver peer relationship.

``-peer-vserver <vserver>`` - Peer Vserver Name::
Specifies name of the peer Vserver with which the Vserver peer
relationship was initiated.

``[-local-name <vserver>]`` - Peer Vserver Local Name::
Specifies the unique local name to identify the peer Vserver with which
the Vserver peer relationship was initiated. The default value is the
remote peer Vserver name.

== Examples

The following example illustrates how to accept the Vserver peer
relationship between Vservers ``_pvs1.example.com_`` residing on
``_cluster2_`` , and ``_lvs1.example.com_`` residing on ``_cluster1_`` .
```

cluster2::> vserver peer accept -vserver pvs1.example.com -peer-vserver lvs1.example.com

```
The following example illustrates how to accept the Vserver peer
relationship between Vservers ``_pvs1.example.com_`` residing on
``_cluster2_`` , and ``_pvs1.example.com_`` residing on ``_cluster1_`` .
During execution of xref:{relative_path}vserver-peer-create.html[vserver
peer create] command on peer cluster, peer Vserver name is locally refered
by unique system generated name ``_pvs1.example.com.1_`` . Using `vserver
peer accept` command specify the unique `-local-name` for peer Vserver.
```

cluster2::> vserver peer accept -vserver pvs1.example.com -peer-vserver pvs1.example.com.1 -local-name locallyUniqueName

```
== Related Links

* xref:{relative_path}vserver-peer-create.html[vserver peer create]




[[IDf1d420a5ba426777599e89e5d85d7680]]
= vserver peer create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new Vserver peer relationship

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer create` command creates a Vserver peer relationship
between two Vservers residing on the same cluster or across two clusters.
For intercluster Vserver peer relationships, the cluster administrator
must accept or reject the relationship on the peer cluster.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Specifies the name of the local Vserver.

``-peer-vserver <vserver>`` - Peer Vserver Name::
```

```
Specifies the name of the peer Vserver with which you want to create the
Vserver peer relationship.

``[-peer-cluster <text>]`` - Peer Cluster Name::
Specifies the name of the peer cluster. If this is not specified, it is
assumed that the peer Vserver resides on the same cluster.

``-applications {snapmirror|file-copy|lun-copy|flexcache}`` - Peering
Applications::
Specifies the applications for which the Vserver peer relationship is
created.

``[-local-name <vserver>]`` - Peer Vserver Local Name::
Specifies the unique local name to identify the peer Vserver with which
you want to create the Vserver peer relationship. The default value is the
remote peer Vserver name.

== Examples

The following example illustrates how to create an intercluster Vserver
peer relationship between Vserver ``_lvs1.example.com_`` , residing on
``_cluster1_`` , and ``_pvs1.example.com_`` , residing on ``_cluster2_`` .
The relationship is created for SnapMirror.
```

cluster1::> vserver peer create -vserver lvs1.example.com -peer-vserver pvs1.example.com -peer-cluster cluster2 -applications snapmirror

```
The following example illustrates how to create an intercluster Vserver
peer relationship between Vserver ``_lvs1.example.com_`` , residing on
``_cluster1_`` , and ``_lvs1.example.com_`` , residing on ``_cluster2_`` .
The relationship is created for SnapMirror. The `-local-name` parameter is
specified to create a local name used to identify the peer Vserver in
cases where the name of the peer Vserver name is not uniquely referenced
from local cluster.
```

cluster1::> vserver peer create -vserver lvs1.example.com -peer-vserver lvs1.example.com -peer-cluster cluster2 -applications snapmirror -local-name cluster2lvs1locallyUniqueName

cluster1::> vserver peer show
Peer Peer Peering Remote
Vserver Vserver State Peer Cluster Applications Vserver
---------- ---------- ---------- ---------- ---------- ---------
lvs1.example.com
cluster2lvs1locallyUniqueName
initiated cluster2 snapmirror lvs1.example.com
cluster1::> vserver peer show -instance
Local Vserver Name: lvs1.example.com

Peer Vserver Name: cluster2lvs1locallyUniqueName
Peering State: initiated
Peering Applications: snapmirror
Remote Vserver Name: lvs1.example.com

```
The following example illustrates how to create an intercluster Vserver
peer relationship between Vserver ``_lvs1_`` , residing on ``_cluster1_``
, and Vserver ``_pvs1_`` , residing on ``_cluster2_`` . The relationship
is created for SnapMirror. The following Vserver peer permission exists on
remote cluster ``_cluster2_`` for local Vserver ``_pvs1_`` .
```

cluster2::> vserver peer permission show
Peer Cluster Vserver Applications
--- --- ---
cluster1 pvs1 snapmirror
1 entries were displayed.

cluster1::> vserver peer create -vserver lvs1 -peer-vserver pvs1 -peer-cluster cluster2 -applications snapmirror

cluster1::> vserver peer show
Peer Peer Peering Remote
Vserver Vserver State Peer Cluster Applications Vserver
--- --- --- --- --- ---
lvs1 pvs1 peered cluster2 snapmirror pvs1

cluster2::> vserver peer show
Peer Peer Peering Remote
Vserver Vserver State Peer Cluster Applications Vserver
--- --- --- --- --- ---
pvs1 lvs1 peered cluster1 snapmirror lvs1

```
Here is an another example which creates an intracluster Vserver peer
relationship.
```

cluster1::> vserver peer create -vserver lvs1.example.com -peer-vserver lvs2.example.com -applications snapmirror

```
[[ID18bd62ff655426c936a433b5b4f02bbe]]
= vserver peer delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Vserver peer relationship

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer delete` command deletes the Vserver peer relationship
between two Vservers.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Specifies the local Vserver name for which you want to delete the Vserver
peer relationship.

``-peer-vserver <vserver>`` - Peer Vserver Name::
Specifies the peer Vserver name with which the Vserver peer relationship
was established.

``[-force <true>]`` - Force Delete::
Deletes the Vserver peer relationship even if the remote cluster is not
accessible due to, for example, network connectivity issues.

``[-foreground {true|false}]`` - Foreground::
This parameter optionally specifies whether the Vserver peer delete
operation can be executed in the background. If nothing is specified, by
default the Vserver peer delete operation is executed in the background.

== Examples

The following example illustrates how to delete the Vserver peer
relationship between two Vservers ``_lvs1.example.com_`` residing on
``_cluster1_`` , and ``_pvs1.example.com_`` residing on ``_cluster2_`` .
```

cluster1::> vserver peer delete -vserver lvs1.example.com -peer-vserver pvs1.example.com

```
[[ID75fa7489386fbae8465469d8b335f7fe]]
= vserver peer modify-local-name
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the local name for a peer Vserver

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer modify-local-name` command modifies the local name for a
remote peer Vserver. The new local name must be unique.

== Parameters

``-peer-cluster <text>`` - Peer Cluster::
Use this parameter to specify the peer cluster.

``-peer-vserver <text>`` - Remote Peer Vserver::
Use this parameter to specify the existing remote peer Vserver name.

``-new-name <vserver>`` - Remote Peer Vserver Local Name::
Use this parameter to specify the new local name of the peer Vserver. The
new local name must conform to the same rules as a Vserver name.

== Examples
```

cluster2::> vserver peer modify-local-name -peer-cluster cluster1 -peer-vserver vs51.example.com -new-name vs51_cluster1.example.com

```
[[ID5d115bbc27278fa3322c378e260228a1]]
= vserver peer modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a Vserver peer relationship

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer modify` command modifies applications of the Vserver
peer relationship.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Specifies name of the local Vserver for which you want to modify
applications of the Vserver peer relationship.

``-peer-vserver <vserver>`` - Peer Vserver Name::
Specifies name of the peer Vserver for which you want to modify
applications of the Vserver peer relationship.

``-applications {snapmirror|file-copy|lun-copy|flexcache}`` - Peering
Applications::
Specifies the Vserver peer applications.

== Examples

The following example illustrates how to modify applications that are part
of the peer relationship between the Vservers ``_lvs1.example.com_``
residing on ``_cluster1_`` , and ``_pvs1.example.com_`` residing on
``_cluster2_`` .
```

cluster1::> vserver peer modify -vserver lvs1.example.com -peer-vserver pvs1.example.com -applications snapmirror

```
[[IDaa4e5aa020aee544453de57b318677d7]]
= vserver peer reject
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reject a Vserver peer relationship

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer reject` command is used to reject the Vserver peer
relationship between the two Vservers. This command is used only for an
intercluster Vserver peer relationship.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Specifies the name of the local Vserver for which you want to reject the
Vserver peer relationship.

``-peer-vserver <vserver>`` - Peer Vserver Name::
Specifies the name of the peer Vserver with which the Vserver peer
relationship was initiated.

== Examples

The following example illustrates how to reject the Vserver peer
relationship between two Vservers ``_lvs1.example.com_`` residing on
``_cluster1_`` , and ``_pvs1.example.com_`` residing on ``_cluster2_`` .
```

cluster1::> vserver peer reject -vserver lvs1.example.com -peer-vserver pvs1.example.com

```
[[ID98b5912778b4db5cad0142db38c779b4]]
= vserver peer repair-peer-name
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Repair the peer vserver name that was not updated during the last rename
operation

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

Updates the peer Vserver name in remote peer clusters for the specified
Vserver in the local cluster.

== Parameters

``-vserver <vserver>`` - vserver [.advaced-priv]#(privilege: advanced)#::
Name of the Vserver in the local cluster. This name will be repaired on
remote peer clusters.

== Examples

The following example updates the peer-Vserver name across the peered
clusters:
```

cluster1::*> vserver peer repair-peer-name -vserver vs1.example.com
Info: Command completed successfully

```
[[ID316784c3def087615cd247f41573b814]]
= vserver peer resume
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resume a Vserver peer relationship

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer resume` command resumes the Vserver peer relationship
between two Vservers.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Specifies name of the local Vserver for which you want to resume the
Vserver peer relationship.

``-peer-vserver <vserver>`` - Peer Vserver Name::
Specifies name of the peer Vserver with which you want to resume the
Vserver peer relationship.

``[-force <true>]`` - Force Resume::
Resumes the Vserver peer relationship even if the remote cluster is not
accessible due to, for example, network connectivity issues.

== Examples

The following example illustrates resuming a Vserver peer relationship
between two Vservers ``_lvs1.example.com_`` residing on ``_cluster1_`` ,
and ``_pvs1.example.com_`` residing on ``_cluster2_`` .
```

cluster1::> vserver peer resume -vserver lvs1.example.com -peer-vserver pvs1.example.com

```
[[ID9e80b6780821067b3d44ddd103bf7633]]
= vserver peer show-all
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
(DEPRECATED)-Display Vserver peer relationships in detail

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer show-all` command displays the following information
about Vserver peer relationships:

*  Local Vserver name
*  Peer Vserver name
*  Local Vserver UUID
*  Peer Vserver UUID
*  Peer cluster name
*  Applications
*  State of the peering relationship
*  Remote Vserver name



== 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.

``[-vserver <vserver>]`` - Local Vserver Name::
If this parameter is specified, the command displays relationships that
match the specified local Vserver.

``[-peer-vserver <text>]`` - Peer Vserver Name::
If this parameter is specified, the command displays relationships that
match the specified peer Vserver.

``[-vserver-uuid <UUID>]`` - Local Vserver UUID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays relationships that
match the specified local Vserver UUID.
```

```
``[-peer-vserver-uuid <UUID>]`` - Peer Vserver UUID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays relationships that
match the specified peer Vserver UUID.

``[-peer-state
{peered|pending|initializing|initiated|rejected|suspended|deleted}]`` -
Peering State::
If this parameter is specified, the command displays relationships that
match the specified peer state.

``[-applications {snapmirror|file-copy|lun-copy|flexcache}]`` - Peering
Applications::
If this parameter is specified, the command displays relationships that
have the specified applications.

``[-peer-cluster <text>]`` - Peer Cluster Name::
If this parameter is specified, the command displays relationships that
have the specified peer cluster name.

``[-remote-vserver-name <text>]`` - Remote Vserver Name::
If this parameter is specified, the command displays relationships that
match the specified remote Vserver.

== Examples

The following example illustrates how to display Vserver peer
relationships.  +
```

cluster1::> vserver peer show-all
Peer Peer Peering Remote
Vserver Vserver State Peer Cluster Applications Vserver
---------- ---------- ---------- ---------- ---------- ---------
lvs1.example.com
lvs2.example.com
peered cluster1 snapmirror lvs2.example.com
lvs1.example.com
pvs1.example.com
peered cluster2 snapmirror pvs1.example.com
lvs2.example.com
lvs1.example.com
peered cluster1 snapmirror lvs1.example.com
lvs3.example.com
pvs1_cluster3.example.com
peered cluster3 snapmirror pvs1.example.com
lvs1.example.com
lvs1_cluster4.example.com
peered cluster4 snapmirror lvs1.example.com

5 entries were displayed.

```
[[ID3f0a73de9efec5d55a86ae48f3f4aced]]
= vserver peer show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Vserver peer relationships

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver peer show` command displays the following information about
Vserver peer relationships:

*  Local Vserver name
*  Peer Vserver name
*  Local Vserver UUID
*  Peer Vserver UUID
*  Peer cluster name
*  State of the peering relationship
*  Applications
*  Remote Vserver name



== 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.

``[-vserver <vserver>]`` - Local Vserver Name::
If this parameter is specified, the command displays relationships that
match the specified local Vserver.
```

```
``[-peer-vserver <text>]`` - Peer Vserver Name::
If this parameter is specified, the command displays relationships that
match the specified peer Vserver.

``[-peer-state
{peered|pending|initializing|initiated|rejected|suspended|deleted}]`` -
Peering State::
If this parameter is specified, the command displays relationships that
match the specified peer state.

``[-applications {snapmirror|file-copy|lun-copy|flexcache}]`` - Peering
Applications::
If this parameter is specified, the command displays relationships that
have the specified applications.

``[-peer-cluster <text>]`` - Peer Cluster Name::
If this parameter is specified, the command displays relationships that
have the specified peer cluster name.

``[-peer-vserver-uuid <UUID>]`` - Peer Vserver UUID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays relationships that
match the specified peer Vserver UUID.

``[-vserver-uuid <UUID>]`` - Local Vserver UUID [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays relationships that
match the specified local Vserver UUID.

``[-remote-vserver-name <text>]`` - Remote Vserver Name::
If this parameter is specified, the command displays relationships that
match the specified remote Vserver.

== Examples

The following examples illustrate how to display Vserver peer
relationships.  + Cluster administrator:
```

cluster1::> vserver peer show
Peer Peer Peering Remote
Vserver Vserver State Peer Cluster Applications Vserver
---------- ---------- ---------- ---------- ---------- ---------
lvs1.example.com
lvs2.example.com
peered cluster1 snapmirror lvs2.example.com
lvs1.example.com
pvs1.example.com

peered cluster2 snapmirror pvs1.example.com
lvs2.example.com
lvs1.example.com
peered cluster1 snapmirror lvs1.example.com
lvs3.example.com
pvs1_cluster3.example.com
peered cluster3 snapmirror pvs1.example.com
lvs1.example.com
lvs1_cluster4.example.com
peered cluster4 snapmirror lvs1.example.com
5 entries were displayed.

```
Vserver administrator:
```

vs11.example.com::> vserver peer show
Peer Peer Peering Remote
Vserver Vserver State Applications Vserver
--- --- --- --- ---
vs11.example.com
pvs21.example.com
peered snapmirror pvs21.example.com
vs11.example.com
vs12.example.com
peered file-copy, snapmirror vs12.example.com
2 entries were displayed.

```
[[IDa88887021b737afc610442fead4ec558]]
= vserver peer suspend
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Suspend a Vserver peer relationship

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer suspend` command suspends the Vserver peer relationship
between two Vservers.

== Parameters

``-vserver <vserver>`` - Vserver Name::
Specifies name of the local Vserver for which you want to suspend the
Vserver peer relationship.

``-peer-vserver <vserver>`` - Peer Vserver Name::
Specifies name of the peer Vserver for which you want to suspend the
Vserver peer relationship.

``[-force <true>]`` - Force Suspend::
Suspends the Vserver peer relationship even if the remote cluster is not
accessible due to, for example, network connectivity issues.

== Examples

The following example illustrates how to suspend the Vserver peer
relationship between two Vservers ``_lvs1.example.com_`` residing on
``_cluster1_`` , and ``_pvs1.example.com_`` residing on ``_cluster2_`` .
```

cluster1::> vserver peer suspend -vserver lvs1.example.com -peer-vserver pvs1.example.com

```
[[ID4ae3db9143ef8fbdf2ff8472b7b2b779]]
= vserver peer permission create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new Vserver peer permission

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer permission create` command creates a new Vserver peer
permission that can be used during intercluster Vserver peer relationship
creation. Once this permission exists for a local Vserver and peer cluster
combination on local cluster, no explicit xref:{relative_path}vserver-
peer-accept.html[vserver peer accept] command is required for any incoming
Vserver peer relationship creation request from a remote cluster for that
local Vserver. Peer relationship directly changes state to ``_peered_`` on
both clusters.

== Parameters

``-peer-cluster <text>`` - Peer Cluster Name::
Specifies the name of the peer Cluster.

``-vserver <text>`` - Vserver Name::
Specifies the name of the local Vserver. Use "*" to create permission that
applies for all local Vservers.

``-applications {snapmirror|flexcache}`` - Peering Applications::
Specifies the applications that can make use of the intercluster Vserver
peer relationship.

== Examples

The following example illustrates how to create Vserver peer permissions:
```

cluster1::> vserver peer permission create -peer-cluster cluster2 -vserver vs1 -applications snapmirror

The following example illustrates how to create a Vserver peer permission that applies for all the local Vservers

cluster1::> vserver peer permission create -peer-cluster cluster2 -vserver "*" -applications snapmirror

```
Warning: This Vserver peer permission applies to all local Vservers. After that no explict
"vserver peer accept" command required for Vserver peer relationship creation request
from peer cluster "cluster2" with any of the local Vservers. Do you
want to continue? {y|n}: y
cluster1::> vserver peer permission show
Peer Cluster Vserver Applications
---------------- --------------- -----------------
cluster2 "*" snapmirror
cluster2 vs1 snapmirror
2 entries were displayed.
```

Note that both all Vservers and any local Vserver name permission can exists at same time.

```
== Related Links

* xref:{relative_path}vserver-peer-accept.html[vserver peer accept]




[[ID5cde5ed2c4608e361a3b2ff4951e9dc2]]
= vserver peer permission delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a Vserver peer permission

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer permission delete` command deletes Vserver peer
permissions.

== Parameters

``-peer-cluster <text>`` - Peer Cluster Name::
Specifies the name of the peer Cluster.

``-vserver <text>`` - Vserver Name::
Specifies the name of the local Vserver.

== Examples

The following example illustrates how to delete Vserver peer permissions:
```

cluster1::> vserver peer permission delete -peer-cluster cluster2 -vserver vs1

```
[[ID0bf32d35274bcd84571ba4b011ca95a0]]
= vserver peer permission modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the Existing Vserver peer permission

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer permission modify` command is used to modify attributes
of the Vserver peer permission relationship. Changes made using this
command will only apply to Vserver peer relationships that are created
after the Vserver peer permission have been modified. Vserver peer
permission is used to give permission to a local Vserver for intercluster
Vserver peer relationship creation so that the command
xref:{relative_path}vserver-peer-accept.html[vserver peer accept] is not
required for incoming Vserver peer relationship creation from a remote
cluster for that local Vserver.

== Parameters

``-peer-cluster <text>`` - Peer Cluster Name::
Specifies the name of the peer cluster.

``-vserver <text>`` - Vserver Name::
Specifies name of the local Vserver for which you want to modify
applications of the Vserver peer permission relationship.

``-applications {snapmirror|flexcache}`` - Peering Applications::
Specifies the applications that can make use of the intercluster Vserver
peer relationship.

== Examples

The following example illustrates how to modify Vserver peer permissions:
```

cluster1::*> vserver peer permission modify -peer-cluster cluster2 -vserver vs1 -applications snapmirror

```
== Related Links
```

```
* xref:{relative_path}vserver-peer-accept.html[vserver peer accept]




[[ID6b08f570687ab81d55ed4300fdede8f1]]
= vserver peer permission show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Vserver peer permissions

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer permission show` command displays the following
information about Vserver peer permissions:

*  Peer cluster name
*  Local Vserver name
*  Applications



== 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.

``[-peer-cluster <text>]`` - Peer Cluster Name::
If this parameter is specified, the command displays permissions that have
the specified peer cluster name.

``[-vserver <text>]`` - Vserver Name::
If this parameter is specified, the command displays permissions that
```

```
match the specified local Vserver.

``[-applications {snapmirror|flexcache}]`` - Peering Applications::
If this parameter is specified, the command displays permissions that have
the specified applications.

== Examples

The following examples illustrate how to display Vserver peer permissions:
```

cluster1::> vserver peer permission show
Peer Cluster Vserver Applications
---------- ---------- ----------
cluster2 "*" snapmirror
cluster3 vs1 snapmirror
2 entries were displayed.

```
[[ID3bfca77b57df38aafeac06f3f45ff797]]
= vserver peer transition create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new transition peer relationship between a 7-Mode system and a
Vserver.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer transition create` command creates a transition peer
relationship between a 7-Mode system and a Vserver. This command is not
supported from 9.12.1

== Parameters

``-local-vserver <vserver name>`` - Local Vserver name::
Specifies the name of the local Vserver.

``-src-filer-name <text>`` - Source 7-Mode system::
Specifies the name of the source 7-Mode system (hostname or IP address).

``[-multi-path-address <text>]`` - Additional address for source 7-Mode
system::
Additional address (hostname or IP address) for the source 7-Mode system.

``[-local-lifs <lif-name>,...]`` - List of Local LIFs::
List of LIFs to be used for this peering relationship. The LIF role can be
data or node-mgmt or intercluster or cluster-mgmt.

== Examples

The following example illustrates how to create a transition peer
relationship between Vserver ``_vs1.example.com_`` , residing on
``_Cluster1_`` , and a 7-Mode system ``_src1.example.com_`` . We can also
specify an additional multipath address ``_src1-e0d.example.com_`` , for
load balancing and list of local LIFs ``_lif1, lif2_`` to be used.
```

Cluster1::> vserver peer transition create -local-vserver vs1.example.com -src-filer-name src1.example.com

-multi-path-address src1-e0d.example.com
-local-lifs lif1,lif2

```
[[ID46371e0b97f7b065eae88953e3295cf0]]
= vserver peer transition delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a transition peer relationship.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer transition delete` command deletes the transition peer
relationship.

== Parameters

``-local-vserver <vserver name>`` - Local Vserver name::
Specifies the name of the local Vserver.

``-src-filer-name <text>`` - Source 7-Mode system::
Specifies the name of the source 7-Mode system(hostname or IP address).

== Examples

The following example illustrates how to delete the transition peer
relationship between a Vserver ``_lvs1.example.com_`` residing on
``_cluster1_`` , and source 7-Mode system ``_src1.example.com_`` .
```

cluster1::> vserver peer transition delete -local-vserver lvs1.example.com -src-filer-name src1.example.com

```
[[ID10a8399b1e0b36529af882460f98ed2a]]
= vserver peer transition modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a transition peer relationship.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

The `vserver peer transition modify` command is used to modify the
multipath address or local LIFs of the transition peer relationship.

== Parameters

``-local-vserver <vserver name>`` - Local Vserver name::
Specifies the name of the local Vserver.

``-src-filer-name <text>`` - Source 7-Mode system::
Specifies the name of the source 7-Mode system (hostname or IP address).

``[-multi-path-address <text>]`` - Additional address for source 7-Mode
system::
Additional address (hostname or IP address) for the source 7-Mode system.

``[-local-lifs <lif-name>,...]`` - List of Local LIFs::
List of LIFs to be used for this peering relationship. The LIF role can be
data or node-mgmt or intercluster or cluster-mgmt.

== Examples

The following example illustrates how to modify a transition peer
relationship`s multipath address.
```

cluster1::> vserver peer transition modify -local-vserver vs1.example.com -src-filer-name src1.example.com -multi-path-address src1-e0b.example.com

```
The following example illustrates how to modify the local LIFs of a
transition peer relationship.
```

Cluster1::> vserver peer transition modify -local-vserver vs1.example.com -src-filer-name src1.example.com -local-lifs lif1,lif2

```
[[ID d4e7bb6d4e7119582b96e738cf945bd2]]
= vserver peer transition show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display transition peer relationships.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver peer transition show` command displays the following
information about transition peer transition relationships:

*  Local Vserver name
*  Source 7-Mode system
*  Multi-path address
*  Local LIFs



== 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.

``[-local-vserver <vserver name>]`` - Local Vserver name::
If this parameter is specified, the command displays transition peer
information about the specified local Vserver.

``[-src-filer-name <text>]`` - Source 7-Mode system::
If this parameter is specified, the command displays transition peer
information about the specified source 7-Mode system.
```

```
``[-multi-path-address <text>]`` - Additional address for source 7-Mode
system::
If this parameter is specified, the command displays information about the
specified multipath-address.

``[-local-lifs <lif-name>,...]`` - List of Local LIFs::
If this parameter is specified, the command displays information about the
specified local LIFs.

== Examples
```

cluster1::> vserver peer transition show
Vserver Source Filer Multi Path Address Local LIFs
------- ------------ ----------------- ----------------
vs1.example.com lif1, lif2
src1.example.com
src1-e0b.example.com

```
:leveloffset: -1


= vserver san commands

:leveloffset: +1


[[IDb9650a2965e8df9c0bfb5d16a97e5d2d]]
= vserver san prepare-to-downgrade
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Restore the SAN Configurations to Earlier Release of Data ONTAP Version.

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command initiates the procedure to restore the configuration to
Earlier Release of ONTAP Version.

As part of this command, capability for making SAN LIF offline if placed
on the DR auxiliary partner as part of LIF placement in MetroCluster
environment will be disabled.

== Parameters

``-feature-set <ClusterVersion>`` - Disable the capability introduced in
the Data ONTAP Version::
Specifies the ONTAP Cluster Version from revert to.

== Examples
```

cluster1::> vserver san prepare-to-downgrade -feature-set 8.3.1

```
Clears the SAN configuration to make it compatible to an earlier ONTAP
release.
```

```
:leveloffset: -1


= vserver security commands

:leveloffset: +1


[[ID1e6e14e0e47a049da72063d9bd2146e7]]
= vserver security file-directory apply
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Apply security descriptors on files and directories defined in a policy to
a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory apply` command applies security
settings to files and directories defined in a security policy of a
Vserver.

Applying a security policy to a Vserver is the last step to creating and
applying NTFS ACLs to files or folders. A security policy contains
definitions for the security configuration of a file (or folder) or set of
files (or, folders). The policy is a container for tasks. A task
associates a file/folder path name to the security descriptor that needs
to be set on the file/folder. Every task in a policy is uniquely
identified by the file/folder path. A policy cannot have duplicate task
entries. There can be only one task per path.

The steps to creating and applying NTFS ACLs are the following:

* Create an NTFS security descriptor.
* Add DACLs and SACLs to the NTFS security descriptor.

NOTE: If you want to audit file and directory events, you must configure
auditing on the Vserver in addition to adding the SACL to the security
descriptor.
```

```
* Create a file/directory security policy.

This step associates the policy with a Vserver.
* Create policy tasks.

A policy task refers to a single operation to apply to a file (or folder)
or to a set of files (or folders). among other things, the task defines
which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.



NOTE: Modifying the ACLs via ONTAP CLI will replace the current
permissions on the directory or path.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver that contains the path to which the security policy
is applied.

``-policy-name <Security policy name>`` - Policy Name::
Specifies the security policy to apply.

``[-ignore-broken-symlinks {true|false}]`` - Skip Broken Symlinks
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter as ``_true_`` , the file-directory apply job
will skip all the symlinks that are broken instead of failing the job.

== Examples

The following example applies a security policy named “p1” to Vserver vs0.
```

cluster1::> vserver security file-directory apply -vserver vs0 -policy-name p1

```
[[ID993a99dab2fa86dbf48f35bc75503b92]]
= vserver security file-directory remove-slag
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Removes Storage-Level Access Guard

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory remove-slag` command removes Storage-
Level Access Guard (SLAG) security from the specified volume or qtree
path.

== Parameters

``-vserver <vserver>`` - Vserver::
Specifies the name of the Vserver that is associated with the volume or
qtree path from where you want to remove SLAG.

``-path <text>`` - Path::
Specifies the volume or qtree mounted junction path from which you want to
remove SLAG security.

== Examples

The following example removes SLAG security from the volume path "/vol1"
on Vserver vs1.
```

cluster1::>vserver security file-directory show -vserver vs1 -path /vol1
Vserver: vs1
File Path: /vol1
Security Style: mixed
Effective Style: unix
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
Unix User Id: 0
Unix Group Id: 0
Unix Mode Bits: 755
Unix Mode Bits in Text: rwxr-xr-x
ACLs: Storage-Level Access Guard security DACL (Applies to Directories):

ALLOW-CIFS1\Administrator-0x1200a9
DACL (Applies to Files):
ALLOW-CIFS1\Administrator-0x1200a9
cluster1::>vserver security file-directory remove-slag -path /vol1 -vserver vs1
cluster1::>vserver security file-directory show -vserver vs1 -path /vol1
Vserver: vs1
File Path: /vol1
Security Style: mixed
Effective Style: unix
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
Unix User Id: 0
Unix Group Id: 0
Unix Mode Bits: 755
Unix Mode Bits in Text: rwxr-xr-x
ACLs: -

```
The following example removes SLAG security from the qtree path "/vol1/q1"
on Vserver vs1.
```

cluster1::>vserver security file-directory show -vserver vs1 -path /vol1/q1
Vserver: vs1
File Path: /vol1/q1
Security Style: mixed
Effective Style: unix
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
Unix User Id: 0
Unix Group Id: 0
Unix Mode Bits: 755
Unix Mode Bits in Text: rwxr-xr-x
ACLs: Storage-Level Access Guard security DACL (Applies to Directories):
ALLOW-CIFS1\Administrator-0x1200a9
DACL (Applies to Files):
ALLOW-CIFS1\Administrator-0x1200a9
cluster1::>vserver security file-directory remove-slag -path /vol1/q1 -vserver vs1
cluster1::>vserver security file-directory show -vserver vs1 -path /vol1/q1
Vserver: vs1
File Path: /vol1/q1
Security Style: mixed
Effective Style: unix
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
Unix User Id: 0
Unix Group Id: 0
Unix Mode Bits: 755
Unix Mode Bits in Text: rwxr-xr-x
ACLs: -

```
[[ID4975163ad43a9b19e01cd26862aa8bff]]
= vserver security file-directory show-effective-permissions
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display effective file or folder permissions

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory show-effective-permissions ` command
displays the effective permission granted to a Windows or UNIX user on the
specified file or folder path. The command output depends on the parameter
or parameters specified with the command.

The `-vserver` , `-win-user-name` or `-unix-user-name` and `-path`
parameters are required for this command. If the optional parameter `-
share-name` is specified, it will display the effective share permission.

== 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.

``-vserver <vserver>`` - Vserver::
Use this required parameter to specify the Vserver that contains the path
to the file or folder specified with the required `-path` parameter. Query
characters, such as "*", are not supported.

{ ``-win-user-name <text>`` - Windows User Name::
Use this parameter to specify the Windows user for which effective
permission needs to be displayed on the given file or folder.

| ``-unix-user-name <text>`` - Unix User Name }::
Use this parameter to specify the UNIX user for which effective permission
needs to be displayed on the given file or folder.
```

```
``-path <text>`` - File Path::
Use this mandatory parameter to specify the path of the file or the folder
for which you want to display effective permissions. The path is relative
to the Vserver root volume. If `-share-name` is specified then the path
will be relative to the share path. Query characters,such as "*", are not
supported.

``[-share-name <Share>]`` - CIFS Share Name::
If you specify this optional parameter, the command displays the file or
directory effective permission for the mentioned user, only for files and
directories contained where the specified path is relative to the root of
the specified share. If this parameter is not specified, the Vserver root
volume is taken as the default. If this optional parameter is specified,
then it will also display the effective share permission of the user.
Wildcard query characters are not supported.

``[-client-ip-address <IP Address>]`` - Client IP Address::
If you specify this optional parameter, the command displays the effective
permission for the user with the specified client ip address.

``[-expand-mask {true|false}]`` - Expand Bit Masks::
If you specify this optional parameter, the command displays effective
permission for files and directories where the hexadecimal bit mask
entries are in expanded bit form. If set to default (false), the command
displays effective permission for file or directory in collapsed (textual)
form.

``[-share-path <text>]`` - CIFS Share Path::
If you specify this parameter, the command displays information only about
the CIFS share that match the specified path. Query characters, such as
"*", are not supported.

``[-permission <Security acl>,...]`` - Effective Permissions::
If you specify this parameter, the command displays effective permission
only if specified permission matches. Wildcard query characters are not
supported.




[[ID35160f3c7a1142a980e13ca9088c906c]]
= vserver security file-directory show
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display file/folder security information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory show` command displays file/folder
security information. The command output depends on the parameter or
parameters specified with the command.

The `-vserver` and `-path` parameters are required for this command. If
you do not specify any of the optional parameters, the command displays
all security information in list format for the specified path.

You can specify the `-fields` parameter to specify which fields of
information to display about files and folders security.

You can specify the `-instance` parameter to display all the security
information in list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``-vserver <vserver>`` - Vserver::
Use this required parameter to specify the Vserver that contains the path
to the file or folder specified with the required `-path` parameter.

{ ``[-path <text>]`` - File Path::
Use this field to specify the path of the file or folder for which you
want to display security information. If the volume name is not specified
in the path, the path is relative to the Vserver root volume. If the
path's last subcomponent has a wildcard ("*"), the output will display
information for all files and directories below the parent path.
+
NOTE: If you want to display information of a file or directory which
contains wildcard ("*") as its last sub-component, then provide the
```

```
complete path inside "<path>".
+

+
For instance, vserver security file-directory show -vserver vs1 -path
"/vol1/*" will show ACL information for the directory named "*", only.

| ``[-inode <integer>]`` - File Inode Number }::
Use this field to specify the inode number of the file or folder for which
you want to display security information. If the volume name is not
specified, inode is searched in the Vserver root volume.

{ ``[-volume-name <volume name>]`` - Volume Name::
If you specify this parameter, the command displays information about file
and directory security only for files and directories where the specified
path is relative to the specified volume. If this parameter is not
specified, the Vserver root volume is taken as default.

| ``[-share-name <Share>]`` - Share Name }::
If you specify this parameter, the command displays information about file
and directory security only for files and directories contained where the
specified path is relative to the root of the specified share. If this
parameter is not specified, the Vserver root volume is taken as default.

``[-lookup-names {true|false}]`` - SID to Name Lookups::
If you specify this parameter, the command displays information about file
and directory security for files and directories where the information
about owner and group are stored as names. If set to false, the command
displays information about file and directory security for files and
directories where the information for owner and group are stored as SIDs.

{ ``[-expand-mask {true|false}]`` - Expand Bit Masks::
If you specify this parameter, the command displays information about file
and directory security for files and directories where the hexadecimal bit
mask entries are in expanded bit form. If set to false, the command
displays information about file and directory security for files and
directories where the hexadecimal bit mask entries are in collapsed form.

| ``[-textual-mask {true|false}]`` - Show Textual Mask::
If you specify this parameter as ``_true_`` , the command displays
information about file and directory security for files and directories
where the hexadecimal bit mask is translated to texual format.

| ``[-sddl {true|false}]`` - Display ACLs in SDDL Format }::
If you specify this parameter, the command displays the ACL information
for files and directories in Security Descriptor Definition Language
```

```
(SDDL) format. If the file has `effective-style` as "unix" then this flag
has no effect.

``[-security-style <security style>]`` - Security Style::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with paths in
volumes of the specified security style.

``[-effective-style <security style>]`` - Effective Style::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
effective security style on the path.

``[-dos-attributes <Hex Integer>]`` - DOS Attributes::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
DOS attributes.

``[-text-dos-attr <TextNoCase>]`` - DOS Attributes in Text::
 If you specify this parameter, the command displays information about
file and directory security only for files and directories with the
specified text DOS attributes.

``[-expanded-dos-attr <TextNoCase>]`` - Expanded Dos Attributes::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
extended DOS attributes. This parameter is useful only for files or
directories where the `-expand-mask` is set to true.

``[-user-id <user name>]`` - UNIX User Id::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
UNIX user ID.

``[-group-id <group name>]`` - UNIX Group Id::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
UNIX group ID.

``[-mode-bits <Octal Permission>]`` - UNIX Mode Bits::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
UNIX mode bits in Octal form.

``[-text-mode-bits <text>]`` - UNIX Mode Bits in Text::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
```

```
UNIX mode bits in text form.

``[-acls <Security acl>,...]`` - ACLs::
If you specify this parameter, the command displays information about file
and directory security only for files and directories with the specified
ACLs. If the specified path is a volume or qtree path and Storage-Level
Access Guard (SLAG) is configured on the volume or qtree, this parameter
displays the SLAG information. It also displays the Dynamic Access Control
(DAC) policies if DAC is configured for the given file or directory path.
The following ACL information can be entered:
+
  * Type of ACL - NTFS or NFSV4
  * Control bits in the security descriptors
  * Owner - only in case of NTFS security descriptors
  * Group - only in case of NTFS security descriptors
  * Access Control Entries - discretionary access control list (DACL) and
system access control list (SACL) access control entries (ACEs) in the ACL



== Examples

The following example displays the security information about the path
"/vol4" in Vserver vs1.
```

cluster1::> vserver security file-directory show -vserver vs1 -path /vol4
(vserver security file-directory show)
Vserver: vs1
File Path: /vol4
File Inode Number: 64
Security Style: ntfs
Effective Style: ntfs
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
Unix User Id: 0
Unix Group Id: 0
Unix Mode Bits: 777
Unix Mode Bits in Text: rwxrwxrwx
ACLs: NTFS Security Descriptor
Control:0x8004
Owner:BUILTIN\Administrators
Group:BUILTIN\Administrators
DACL - ACEs
ALLOW-Everyone-0x1f01ff
ALLOW-Everyone-0x10000000-OI|CI|IO

```
The following example displays the security information about the path
"/a/b/file.txt" in Vserver vs1.
```

cluster1::> vserver security file-directory show -vserver vs1 -path /a/b/file.txt -volume-name vol1
(vserver security file-directory show)
Vserver: vs1
File Path: /vol1/a/b/file.txt
File Inode Number: 101
Security Style: ntfs
Effective Style: ntfs
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
Unix User Id: 0
Unix Group Id: 0
Unix Mode Bits: 777
Unix Mode Bits in Text: rwxrwxrwx
ACLs: NTFS Security Descriptor
Control:0x8004
Owner:BUILTIN\Administrators
Group:BUILTIN\Administrators
DACL - ACEs
ALLOW-Everyone-0x1f01ff
ALLOW-Everyone-0x10000000-OI|CI|IO

```
The following example displays the security information of the volume path
"/vol1" containing SLAG.
```

cluster1::> vserver security file-directory show -vserver vs1 -path /vol1
Vserver: vs1
File Path: /vol1
File Inode Number: 64
Security Style: mixed
Effective Style: ntfs
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attribute: -
Unix User Id: 0
Unix Group Id: 1
Unix Mode Bits: 777
Unix Mode Bits in Text: rwxrwxrwx
ACLs: NTFS Security Descriptor
Control:0xbf14
Owner:CIFS1\Administrator
Group:CIFS1\Domain Admins
SACL - ACEs
ALL-Everyone-0xf01ff-OI|CI|SA|FA
RESOURCE ATTRIBUTE-Everyone-0x0
("Department_MS",TS,0x10020,"Finance")
POLICY ID-All resources - No Write-0x0-OI|CI

DACL - ACEs
ALLOW-CIFS1\Administrator-0x1f01ff-OI|CI
ALLOW-Everyone-0x1f01ff-OI|CI
ALLOW CALLBACK-DAC\skanyal-0x1200a9-OI|CI
((@User.department==@Resource.Department_MS@Resource.Impact_MS>1000)@Device.department==@Resource.Department_MS)
Storage-Level Access Guard security
SACL (Applies to Directories):
AUDIT-R1\user1-0x001f01ff-FA
DACL (Applies to Directories):
ALLOW-R1\user1-0x001f01ff
ALLOW-R1\user2-0x001200a9
SACL (Applies to Files):
AUDIT-R1\user1-0x001f01ff-FA
DACL (Applies to Files):
ALLOW-R1\user1-0x001f01ff
ALLOW-R1\user2-0x001200a9

```
The following example displays the security information of the qtree path
"/vol1/q1" containing SLAG.
```

cluster1::> vserver security file-directory show -vserver vs1 -path /vol1/q1
Vserver: vs1
File Path: /vol1/q1
File Inode Number: 105
Security Style: mixed
Effective Style: ntfs
DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attribute: -
Unix User Id: 0
Unix Group Id: 1
Unix Mode Bits: 777
Unix Mode Bits in Text: rwxrwxrwx
ACLs: NTFS Security Descriptor
Control:0xbf14
Owner:CIFS1\Administrator
Group:CIFS1\Domain Admins
SACL - ACEs
ALL-Everyone-0xf01ff-OI|CI|SA|FA
DACL - ACEs
ALLOW-CIFS1\Administrator-0x1f01ff-OI|CI
ALLOW-Everyone-0x1f01ff-OI|CI
Storage-Level Access Guard security
SACL (Applies to Directories):
AUDIT-R1\user1-0x001f01ff-FA
DACL (Applies to Directories):
ALLOW-R1\user1-0x001f01ff
ALLOW-R1\user2-0x001200a9
SACL (Applies to Files):
AUDIT-R1\user1-0x001f01ff-FA
DACL (Applies to Files):

ALLOW-R1\user1-0x001f01ff
ALLOW-R1\user2-0x001200a9

```
[[ID51051b7bfe3ee771a9adbf53e704bcdc]]
= vserver security file-directory job show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a list of file security jobs

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory job show` command displays
information about security file-directory jobs.

To display detailed information about a specific job, run the command with
the `-id` parameter.

You can specify additional parameters to select information that matches
the values you specify for those parameters. For example, to display
information only about security file-directory jobs running on a specific
node, run the command with the `-node` parameter.

== 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.

| ``[-inprogress  ]``::
Displays the job ID, the job name, the owning Vserver, and the progress of
the security file-directory job.

| ``[-jobstate  ]``::
Displays information about each job's state, including the queue state,
whether the job was restarted and when the job has completely timed out.

| ``[-jobuuid  ]`` [.advaced-priv]#(privilege: advanced)#::
Displays the job ID, the job name, the owning Vserver, and the job UUID.
```

```
| ``[-sched  ]``::
Displays the job ID, the job name, the owning Vserver, and the schedule on
which the security file-directory job runs.

| ``[-times  ]``::
Displays the job ID, the job name, the owning Vserver, the time when the
job was last queued, the time when the job was last started, and the time
when the job most recently ended.

| ``[-type  ]``::
Displays the job ID, the job name, the job type, and the job category.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-id <integer>]`` - Job ID::
Selects the jobs that match the ID or range of IDs that you specify.

``[-vserver <vserver name>]`` - Owning Vserver::
Selects jobs that are owned by the specified Vserver.

``[-name <text>]`` - Name::
Selects the jobs that match this parameter value.

``[-description <text>]`` - Description::
Selects the jobs that match this parameter value.

``[-priority {Low|Medium|High|Exclusive}]`` - Priority::
Selects the jobs that match this parameter value.

``[-node <nodename>]`` - Node::
Selects the jobs that match this parameter value.

``[-affinity {Cluster|Node}]`` - Affinity::
Selects the jobs that match this parameter value.

``[-schedule <job_schedule>]`` - Schedule::
Selects the jobs that match this parameter value.

``[-queuetime <MM/DD HH:MM:SS>]`` - Queue Time::
Selects the jobs that match this parameter value.

``[-starttime <MM/DD HH:MM:SS>]`` - Start Time::
Selects the jobs that match this parameter value.
```

```
``[-endtime <MM/DD HH:MM:SS>]`` - End Time::
Selects the jobs that match this parameter value.

``[-dropdeadtime <MM/DD HH:MM:SS>]`` - Drop-dead Time::
Selects the jobs that match this parameter value.

``[-restarted {true|false}]`` - Restarted?::
Selects the jobs that match this parameter value.

``[-state
{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Re
schedule|Error|Quit|Dead|Unknown|Restart|Dormant}]`` - State::
Selects the jobs that match this parameter value.

``[-code <integer>]`` - Status Code::
Selects the jobs that match this parameter value.

``[-completion <text>]`` - Completion String::
Selects the jobs that match this parameter value.

``[-jobtype <text>]`` - Job Type::
Selects the jobs that match this parameter value.

``[-category <text>]`` - Job Category::
Selects the jobs that match this parameter value.

``[-uuid <UUID>]`` - UUID::
Selects the jobs that match this parameter value.

``[-progress <text>]`` - Execution Progress::
Selects the jobs that match this parameter value.

``[-username <text>]`` - User Name::
Selects the jobs that match this parameter value.

``[-process <text>]`` - Process::
Selects jobs with the specified process number.

== Examples

The following example displays information bathe file-directory security
job.
```

cluster1::> vserver security file-directory apply -policy-name pol -vserver vs1
cluster1::> vserver security file-directory job show
Owning
Job ID Name Vserver Node State

--- --- --- --- ---
25 Fsecurity Apply vsim2.3 vsim2.3-01 Success
Description: File Directory Security Apply Job

```
[[ID87f5ce96cb098917fe19fd585fbdfd6e]]
= vserver security file-directory ntfs create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an NTFS security descriptor

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs create` command creates an NTFS
security descriptor to which you can add access control entries (ACEs) to
the discretionary access control list (DACL) and the system access control
list (SACL).

Creating an NTFS security descriptor is the first step in configuring and
applying NTFS access control lists (ACLs) to files and folders residing
within a namespace. Later, you will associate the security descriptor to a
policy task.

You can create NTFS security descriptors for files and folders residing
within FlexVol volumes with NTFS security-style or on NTFS security
descriptors on mixed security-style volumes.

The steps to creating and applying NTFS ACLs are the following:

* Create an NTFS security descriptor.
* Add DACLs and SACLs to the NTFS security descriptor.

NOTE: If you want to audit file and directory events, you must configure
auditing on the Vserver in addition to adding a SACL to the security
descriptor.


* Create a file/directory security policy.

This step associates the policy with a Vserver.
```

```
* Create a policy task.

A policy task refers to a single operation to apply to a file (or folder)
or to a set of files (or folders). among other things, the task defines
which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver on which to create the security
descriptor.

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor you want to create. After
you create a security descriptor, you can add SACL and DACL access control
entries (ACEs) to it.
+
NOTE:  Every newly created security descriptor contains the 4 default DACL
ACEs as mentioned below:
```

Vserver: vserver1
NTFS Security Descriptor Name: sd1
Account Name Access Access Apply To
Type Rights
--- --- --- ---
BUILTIN\Administrators
allow full-control this-folder, sub-folders, files
BUILTIN\Users allow full-control this-folder, sub-folders, files
CREATOR OWNER allow full-control this-folder, sub-folders, files
NT AUTHORITY\SYSTEM
allow full-control this-folder, sub-folders, files

```
+



``[-owner <name or sid>]`` - Owner::
Specifies the owner of the security descriptor. You can specify the owner
using either a user name or SID.
+
The owner of the security descriptor can modify the permissions on the
file (or folder) or files (or folders) to which the security descriptor is
applied and can give other users the right to take ownership of the object
or objects to which the security descriptor is applied. You can use any of
the following formats when specifying the value for this parameter:
```

```
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN


+
NOTE: If you specify any of the three user name formats for the value of
`-owner` , keep in mind that the value for the user name is case
insensitive. The value for the user name is ignored for Storage-Level
Access Guard (SLAG).
+


``[-group <name or sid>]`` - Primary Group [.advaced-priv]#(privilege:
advanced)#::
Specifies the owner group of the security descriptor. You can specify the
owner group using either a group name or SID. You can use any of the
following formats when specifying the value for this parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN


+
NOTE: If you specify any of the three user name formats for the value of
`-group` , keep in mind that the value for the user name is case
insensitive. The value for the user name is ignored for SLAG.
+


``[-control-flags-raw <Hex Integer>]`` - Raw Control Flags [.advaced-
priv]#(privilege: advanced)#::
Specifies the control flags in the security descriptor.
+
NOTE: The value for the control flag is ignored for SLAG.
+


== Examples

The following example creates an NTFS security descriptor named “sd1” on
```

```
Vserver "vs1" and assigns “DOMAIN\Administrator” as the security
descriptor owner.
```

cluster1::> vserver security file-directory ntfs create -ntfs-sd sd1 -vserver vs1 -owner DOMAIN\Administrator
cluster1::> vserver security file-directory ntfs show -vserver vs1 -ntfs-sd sd1
Vserver: vs1
Security Descriptor Name: sd2
Owner of the Security Descriptor: DOMAIN\Administrator

```
[[ID9a20acf3dfcbb88351a897f03a1b474f]]
= vserver security file-directory ntfs delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an NTFS security descriptor

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs delete` command deletes an NTFS
security descriptor. Deleting a security descriptor also deletes all the
contained DACL and SACL access control entries (ACEs).

== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver that is associated with the security
descriptor that you want to delete.

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor to delete.

== Examples

The following example deletes an NTFS security descriptor named "sd1" on
Vserver vs1.
```

cluster1::> vserver security file-directory ntfs delete -ntfs-sd sd1 -vserver vs1

```
[[ID8e85258d2e5ef146d0f1dad2c7032abb]]
```

```
= vserver security file-directory ntfs modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an NTFS security descriptor

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs modify` command modifies an NTFS
security descriptor. You can change the `-owner` , `-group` and `-control-
flags-raw ` of the security descriptor with this command.

== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security descriptor
that you want to modify.

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor that you want to modify.

``[-owner <name or sid>]`` - Owner::
Specifies the owner of the security descriptor. You can specify the owner
using either the user name or SID.
+
The owner of the security descriptor can modify the permissions on the
file (or folder) or files (or folders) to which the security descriptor is
applied and can give other users the right to take ownership of the object
or objects to which the security descriptor is applied. You can use any of
the following formats when specifying the value for this parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN



+
NOTE: If you specify any of the three user name formats for the value of
`-owner` , keep in mind that the value for the user name is case
```

```
insensitive. The value for the user name is ignored for Storage-Level
Access Guard (SLAG).
+


``[-group <name or sid>]`` - Primary Group [.advaced-priv]#(privilege:
advanced)#::
Specifies the owner group of the security descriptor. You can specify the
owner group using either a group name or SID. You can use any of the
following formats when specifying the value for this parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN



+
NOTE: If you specify any of the three user name formats for the value of
`-group` , keep in mind that the value for the user name is case
insensitive. The value for the user name is ignored for SLAG.
+


``[-control-flags-raw <Hex Integer>]`` - Raw Control Flags [.advaced-
priv]#(privilege: advanced)#::
Specifies the control flags in the security descriptor to be modified.
+
NOTE: The value for the control flag is ignored for SLAG.
+


== Examples

The following example modifies the owner of an NTFS security descriptor
named "sd2" on Vserver vs1.
```

cluster1::> vserver security file-directory ntfs modify -ntfs-sd sd2 -vserver vs1 -owner domain\administrator
cluster1::> vserver security file-directory ntfs show -vserver vs1 -ntfs-sd sd2
Vserver: vs1
Security Descriptor Name: sd2
Owner of the Security Descriptor: DOMAIN\Administrator

```
[[ID8387717bed224d06a240d7d7ab3ae1e1]]
= vserver security file-directory ntfs show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display an NTFS security descriptors

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs show` command displays
information about the security descriptor. The command output depends on
the parameter or parameters specified with the command. If you do not
specify any parameters, the command displays all information about all
security descriptors defined on the cluster.

You can specify the `-fields` parameter to specify which fields of
information to display about security descriptors.

You can specify the `-instance` parameter to display all the information
about security descriptors in list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the security descriptors associated with the Vserver that you specify.

``[-ntfs-sd <ntfs sd name>]`` - NTFS Security Descriptor Name::
If you specify this parameter, the command displays information only about
the security descriptors that you specify.

``[-owner <name or sid>]`` - Owner::
If you specify this parameter, the command displays information only about
the security descriptors owned by the specified user name or SID.
```

```
``[-group <name or sid>]`` - Primary Group [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the security descriptors associated with the owner group.

``[-control-flags-raw <Hex Integer>]`` - Raw Control Flags [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the security descriptors associated with the control flags.

== Examples

The following example displays information about an NTFS security
descriptor named “sd2” on Vserver vs1.
```

cluster1::> vserver security file-directory ntfs show -vserver vs1 -ntfs-sd sd2
Vserver: vs1
Security Descriptor Name: sd2
Owner of the Security Descriptor: DOMAIN\Administrator

```
[[ID0466619685e3f9c48ab26c89fe5f00a5]]
= vserver security file-directory ntfs dacl add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a DACL entry to NTFS security descriptor

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs dacl add` command adds access
control entries (ACEs) into a security descriptor’s discretionary access
control list (DACL).

If the security descriptor contains a DACL that has existing ACEs, the
command adds the new ACE to the DACL. If the security descriptor does not
contain a DACL, the command creates the DACL and adds the new ACE to it.

Adding a DACL entry to the security descriptor is the second step in
configuring and applying ACLs to a file or folder. Before you can add a
DACL entry to a security descriptor, you must first create the security
```

```
descriptor.

The steps to creating and applying NTFS ACLs are the following:

* Create an NTFS security descriptor.
* Add DACLs and SACLs to the NTFS security descriptor.

NOTE: If you want to audit file and directory events, you must configure
auditing on the Vserver in addition to adding the SACL to the security
descriptor.


* Create a file/directory security policy.

This step associates the policy with a Vserver.
* Create policy tasks.

A policy task refers to a single operation to apply to a file (or folder)
or to a set of files (or folders). among other things, the task defines
which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security descriptor
to which you want to add a discretionary access control entry
(discretionary ACE).

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor to which you want to add a
discretionary access control entry.

``-access-type {deny|allow}`` - Allow or Deny::
Specifies whether the discretionary access control entry is an ``_allow_``
or ``_deny_`` type of access control.

``-account <name or sid>`` - Account Name or SID::
Specifies the account on which to apply the discretionary access control
entry. You can specify the account by using a user name or SID. You can
use any of the following formats when specifying the value for this
parameter:
+
+
  * SID
```

```
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN


+
NOTE: If you specify any of the three user name formats for the value of
-account, keep in mind that the value for the user name is case
insensitive.
+


{ ``[-rights {no-access|full-control|modify|read-and-
execute|read|write}]`` - DACL ACE's Access Rights::
Specifies the right that you want to add for the account specified in the
`-account` parameter. The `-rights` parameter is mutually exclusive with
the `-advanced-rights` and `-rights-raw` parameter. If you specify the `-
rights` parameter, you can only specify one value.
+
You can specify one of the following rights values:
+
  * no-access
  * full-control
  * modify
  * read-and-execute
  * read
  * write



| ``[-advanced-rights <Advanced access right>,...]`` - DACL ACE's Advanced
Access Rights }::
Specifies the advanced rights that you want to add for the account
specified in the `-account` parameter. The `-advanced-rights` parameter is
mutually exclusive with the `-rights` and `-rights-raw` parameter. You can
specify more than one advanced-rights value by using a comma-delimited
list.
+
You can specify one or more of the following advanced rights:
+
  * read-data
  * write-data
  * append-data
  * read-ea
  * write-ea
  * execute-file
```

```
  * delete-child
  * read-attr
  * write-attr
  * delete
  * read-perm
  * write-perm
  * write-owner
  * full-control



| ``[-rights-raw <Hex Integer>]`` - DACL ACE's Raw Access Rights
[.advaced-priv]#(privilege: advanced)# }::
Specifies the raw rights that you want to add for the account specified in
the `-account` parameter. The `rights-raw` parameter is mutually exclusive
with the `-advanced-rights` and `-rights` parameter. Specify the value as
a hexadecimal integer, for example: ``_0xA10F_`` or ``_0xb3ff_`` etc.

``[-apply-to {this-folder|sub-folders|files}]`` - Apply DACL Entry::
Specifies where to apply the discretionary access control entry. You can
specify more than one value by using a comma-delimited list.
+
You can specify one or more of the following values:
+
  * this-folder
  * sub-folder
  * files



+
NOTE: Select one of the following combinations of values for the `-apply-
to` parameter for Storage-Level Access Guard (SLAG):
+
  * this-folder, sub-folder, files
  * this-folder, sub-folder
  * files



+
If you specify an invalid `-apply-to` value, this security descriptor is
removed from the associated Storage-Level Access Guard (SLAG) `security
file-directory policy task` .
+


== Examples
```

```
The following example adds a DACL entry to the security descriptor named
“sd1” on Vserver "vs1" for the "DOMAIN\Administrator" account.
```

cluster1::> vserver security file-directory ntfs dacl add -ntfs-sd sd1 -access-type deny -account DOMAIN\Administrator -rights full-control -apply-to this-folder -vserver vs1
cluster1::> vserver security file-directory ntfs dacl show -vserver vs1 -ntfs-sd sd1 -access-type deny -account domain\administrator
Vserver: vs1
Security Descriptor Name: sd1
Allow or Deny: deny
Account Name or SID: DOMAIN\Administrator
Access Rights: full-control
Advanced Access Rights: -
Apply To: this-folder
Access Rights: full-control

```
[[ID4852aa01e60537747c3ed90e6e2c843b]]
= vserver security file-directory ntfs dacl modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an NTFS security descriptor DACL entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs dacl modify` command modifies
parameters in an existing discretionary access control (DACL) entry.

You can unambiguously define which DACL entry to modify by specifying the
following four parameters in the modify command:

* Vserver associated with the security descriptor that contains the DACL
entry
* Name of the security descriptor that contains the DACL entry
* Whether the DACL is an allow or deny type of DACL entry
* The account name or SID to which the DACL is applied



You can modify the following parameters:
```

```
* -right,-advanced-rights ,-rights-raw
* -apply-to



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security descriptor
containing the discretionary access control entry whose parameters you
want to modify.

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor that contains the
discretionary access control entry that you want to modify.

``-access-type {deny|allow}`` - Allow or Deny::
Specifies whether the discretionary access control entry that you want to
modify is an ``_allow_`` or ``_deny_`` type of access control.

``-account <name or sid>`` - Account Name or SID::
Specifies the account associated with the discretionary access control
entry you want to modify. You can specify the account by using a user name
or SID. You can use any of the following formats when specifying the value
for this parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN



+
NOTE: If you specify any of the three user name formats for the value of
-account, keep in mind that the value for the user name is case
insensitive.
+



{ ``[-rights {no-access|full-control|modify|read-and-
execute|read|write}]`` - Access Rights::
Specifies the right that you want to add for the account specified in the
`-account` parameter. The `-rights` parameter is mutually exclusive with
the `-advanced-rights` and `-rights-raw` parameter. If you specify the `-
rights` parameter, you can only specify one value.
+
```

```
You can specify one of the following rights values:
+
  * no-access
  * full-control
  * modify
  * read-and-execute
  * read
  * write



| ``[-rights-raw <Hex Integer>]`` - Raw Access Rights [.advaced-
priv]#(privilege: advanced)#::
Specifies the raw rights that you want to add for the account specified in
the `-account` parameter. The `-rights-raw` parameter is mutually
exclusive with the `-advanced-rights` and `-rights` parameter. Specify the
value as a hexadecimal integer, for example: ``_0xA10F_`` or ``_0xb3ff_``
etc.

| ``[-advanced-rights <Advanced access right>,...]`` - Advanced Access
Rights }::
Specifies the advanced rights that you want to add for the account
specified in the `-account` parameter. The `-advanced-rights` parameter is
mutually exclusive with the `-rights` and `-rights-raw` parameter. You can
specify more than one advanced-rights value by using a comma-delimited
list.
+
You can specify one or more of the following advanced rights:
+
  * read-data
  * write-data
  * append-data
  * read-ea
  * write-ea
  * execute-file
  * delete-child
  * read-attr
  * write-attr
  * delete
  * read-perm
  * write-perm
  * write-owner
  * full-control
```

```
``[-apply-to {this-folder|sub-folders|files}]`` - Apply DACL Entry::
Specifies where to apply the discretionary access control entry. You can
specify more than one value by using a comma-delimited list.
+
You can specify one or more of the following values:
+
  * this-folder
  * sub-folder
  * files



+
NOTE: Select one of the following combinations of values for the `-apply-
to` parameter for Storage-Level Access Guard (SLAG):
+
  * this-folder, sub-folder, files
  * this-folder, sub-folder
  * files



+
If you specify an invalid `-apply-to` value, this security descriptor is
removed from the associated Storage-Level Access Guard (SLAG) `security
file-directory policy task` .
+


== Examples

The following example modifies the `-right` and `-apply-to` parameters in
the DACL entry associated to the security descriptor named “sd2” on
Vserver vs1 for the "BUILTIN\Administrators" account.
```

cluster1::> vserver security file-directory ntfs dacl modify -ntfs-sd sd2 -access-type allow -account BUILTIN\Administrators -vserver vs1 -rights modify -apply-to this-folder,sub-folders
cluster1::> vserver security file-directory ntfs dacl show -vserver vs1 -ntfs-sd sd2 -account BUILTIN\Administrators -instance
Vserver: vs1
Security Descriptor Name: sd2
Allow or Deny: allow
Account Name or SID: BUILTIN\Administrators
Access Rights: modify
Advanced Access Rights: -
Apply To: this-folder, sub-folders
Access Rights: modify

```
[[IDf7546c83811f20888203bd3ccfb282fb]]
```

```
= vserver security file-directory ntfs dacl remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a DACL entry from NTFS security descriptor.

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs dacl remove` command removes a
discretionary access control entry from a security descriptor.

You can unambiguously define which DACL entry to remove by specifying the
following four parameters in the command:

* Vserver associated with the security descriptor that contains the DACL
entry
* Name of the security descriptor that contains the DACL entry
* Whether the DACL is an allow or deny type of DACL entry
* The account name or SID to which the DACL is applied



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security descriptor
from which you want to remove a discretionary access control entry.

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor that contains the
discretionary access control entry that you want to remove.

``-access-type {deny|allow}`` - Allow or Deny::
Specifies whether the discretionary access control entry you want to
remove is an ``_allow_`` or ``_deny_`` of access control.

``-account <name or sid>`` - Account Name or SID::
Specifies the account name or SID associated with the discretionary access
control entry that you want to remove.

== Examples
```

```
The following example removes a DACL entry from the security descriptor
named “sd2” with “allow” access type for the "BUILTIN\Administrators"
account on Vserver vs1.
```

cluster1::> vserver security file-directory ntfs dacl remove -ntfs-sd sd2 -access-type allow -account BUILTIN\Administrators -vserver vs1

```
[[ID6418b73ea2628cb8c4dec79ba3f78799]]
= vserver security file-directory ntfs dacl show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NTFS security descriptor DACL entries

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs dacl show` command displays
information about all the discretionary access control entries in the
Vserver. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information about all DACL entries:

* Vserver name
* Security descriptor
* List of DACL entries



You can specify the `-fields` parameter to specify which fields of
information to display about DACL entries.

You can specify the `-instance` parameter to display all information about
DACL entries in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.
```

```
| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
discretionary access control entries associated with the specified
Vserver.

``[-ntfs-sd <ntfs sd name>]`` - NTFS Security Descriptor Name::
If you specify this parameter, the command displays information only about
the discretionary access control entries for the security descriptor that
you specify.

``[-access-type {deny|allow}]`` - Allow or Deny::
If you specify this parameter, the command displays information only about
the discretionary access control entries with the access type that you
specify.

``[-account <name or sid>]`` - Account Name or SID::
If you specify this parameter, the command displays information only about
the discretionary access control entries associated with the account name
or SID that you specify. You can use any of the following formats when
specifying the value for this parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN



+
NOTE: If you specify any of the three user name formats for the value of
-account, keep in mind that the value for the user name is case
insensitive.
+



``[-rights {no-access|full-control|modify|read-and-execute|read|write}]``
- Access Rights::
If you specify this parameter, the command displays information only about
the discretionary access control entries with the user right that you
specify. Only one value can be specified.
+
```

```
You can specify one of the following rights values:
+
  * no-access
  * full-control
  * modify
  * read-and-execute
  * read
  * write



``[-rights-raw <Hex Integer>]`` - Raw Access Rights [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the discretionary access control entries with the advanced user rights
that you specify. This value for this parameter is mutually exclusive with
any other rights values. Specify the value as a hexadecimal integer, for
example: ``_0xA10F_`` or ``_0xb3ff_`` etc.

``[-advanced-rights <Advanced access right>,...]`` - Advanced Access
Rights::
If you specify this parameter, the command displays information only about
the discretionary access control entries with the advanced user rights
that you specify. You can specify more than one value by using a comma-
delimited list.
+
You can specify one or more of the following advanced rights:
+
  * read-data
  * write-data
  * append-data
  * read-ea
  * write-ea
  * execute-file
  * delete-child
  * read-attr
  * write-attr
  * delete
  * read-perm
  * write-perm
  * write-owner
  * full-control



``[-apply-to {this-folder|sub-folders|files}]`` - Apply DACL Entry::
```

```
If you specify this parameter, the command displays information only about
the discretionary access control entries with the -applied-to value or
values that you specify. You can specify more than one value by using a
comma-delimited list.
+
You can specify one or more of the following values:
+
  * this-folder
  * sub-folder
  * files



``[-readable-access-rights <TextNoCase>]`` - Access Rights::
If you specify this parameter, the command displays information only the
discretionary access control entries with the readable access rights that
you specify.

== Examples

The following example shows information about a DACL entry.
```

cluster1::> vserver security file-directory ntfs dacl show
Vserver: vs1
NTFS Security Descriptor Name: sd2
Account Name Access Access Apply To
Type Rights
---------- ------- ------- ----------
BUILTIN\Users allow full-control this-folder, sub-folders, files
CREATOR OWNER allow full-control this-folder, sub-folders, files
NT AUTHORITY\SYSTEM
allow full-control this-folder, sub-folders, files
3 entries were displayed.

```
[[ID95d0abd4173b520f57044a2ee30f0e94]]
= vserver security file-directory ntfs sacl add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a SACL entry to NTFS security descriptor

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver security file-directory ntfs sacl add` command adds system
access control list entries (ACEs) into a security descriptor's system
access control list (SACL).

If the security descriptor contains a SACL that has existing security
ACEs, the command adds the new security ACE to the SACL. If the security
descriptor does not contain a SACL, the command creates the SACL and adds
the new security ACE to it.

Adding a SACL entry to the security descriptor is the second step in
configuring and applying security ACLs to a file or folder. Before you can
add a SACL entry to a security descriptor, you must first create the
security descriptor.

The steps to creating and applying NTFS ACLs are the following:

* Create an NTFS security descriptor.
* Add DACL and SACL entries to the NTFS security descriptor.

NOTE: If you want to audit file and directory events, you must configure
auditing on the Vserver in addition to adding the SACL to the security
descriptor.


* Create a file/directory security policy.

This step associates the policy with a Vserver.
* Create policy tasks.

A policy task refers to a single operation to apply to a file (or folder)
or to a set of files (or folders). among other things, the task defines
which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security descriptor
to which you want to add a system access control list entry.

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor to which you want to add a
system access control list entry.
```

```
``-access-type {failure|success}`` - Success or Failure::
Specifies whether the system access control list entry that you want to
add is a ``_failure_`` or ``_success_`` access audit type.

``-account <name or sid>`` - Account Name or SID::
Specifies the account on which to apply the system access control list
entry. You can specify the account by using a user name or SID. You can
use any of the following formats when specifying the value for this
parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN


+
NOTE: If you specify any of the three user name formats for the value of
`-account` , keep in mind that the value for the user name is case
insensitive.
+


{ ``[-rights {no-access|full-control|modify|read-and-
execute|read|write}]`` - Access Rights::
Specifies the rights that you want to get audited for the account
specified in the `-account` parameter. The `-rights` parameter is mutually
exclusive with the `-advanced-rights` and `-rights-raw` parameter. If you
specify the `-rights` parameter, you can only specify one value.
+
You can specify one of the following rights values:
+
  * no-access
  * full-control
  * modify
  * read-and-execute
  * read
  * write



| ``[-advanced-rights <Advanced access right>,...]`` - Advanced Access
Rights }::
Specifies the advanced rights that you want to get audited for the account
```

```
specified in the `-account` parameter. The `-advanced-rights` parameter is
mutually exclusive with the `-rights` and `-rights-raw` parameter. You can
specify more than one advanced-rights value by using a comma-delimited
list.
+
You can specify one or more of the following advanced rights:
+
  * read-data
  * write-data
  * append-data
  * read-ea
  * write-ea
  * execute-file
  * delete-child
  * read-attr
  * write-attr
  * delete
  * read-perm
  * write-perm
  * write-owner
  * full-control



| ``[-rights-raw <Hex Integer>]`` - Raw Access Rights [.advaced-
priv]#(privilege: advanced)# }::
Specifies the raw rights that you want to get audited for the account
specified in the `-account` parameter. The `-rights-raw` parameter is
mutually exclusive with the `-advanced-rights` and `-rights` parameter.
Specify the value as a hexadecimal integer, for example: ``_0xA10F_`` or
``_0xb3ff_`` etc.

``[-apply-to {this-folder|sub-folders|files}]`` - Apply SACL To::
Specifies where to apply the system access control list entry. You can
specify more than one value by using a comma-delimited list.
+
You can specify one or more of the following values:
+
  * this-folder
  * sub-folder
  * files



+
NOTE: Select one of the following combinations of values for the `-apply-
to` parameter for Storage-Level Access Guard (SLAG):
```

```
+
  * this-folder, sub-folder, files
  * this-folder, sub-folder
  * files


+
If you specify an invalid `-apply-to` value, this security descriptor is
removed from the associated Storage-Level Access Guard (SLAG) `security
file-directory policy task` .
+


== Examples

The following example adds a SACL entry to the security descriptor named
“sd1” on Vserver vs1.
```

cluster1::> vserver security file-directory ntfs sacl add -ntfs-sd sd1 -access-type failure -account DOMAIN\Administrator -rights full-control -apply-to this-folder -vserver vs1
cluster1::> vserver security file-directory ntfs sacl show -vserver vs1 -ntfs-sd sd1 -access-type deny -account DOMAIN\Administrator
Vserver: vs1
Security Descriptor Name: sd1
Access type for Specified Access Rights: failure
Account Name or SID: DOMAIN\Administrator
Access Rights: full-control
Advanced Access Rights: -
Apply To: this-folder
Access Rights: full-control

```
[[ID891b940f80c8ae84d1adaf4f68bbf010]]
= vserver security file-directory ntfs sacl modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an NTFS security descriptor SACL entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs sacl modify` command modifies
```

```
parameters in an existing system access control list entry.

You can unambiguously define which SACL entry to modify by specifying the
following four parameters in the modify command:

* Vserver associated with the security descriptor that contains the SACL
entry
* Name of the security descriptor that contains the SACL entry
* Whether the SACL is a success or failure type of SACL entry
* The account name or SID to which the SACL is applied



You can modify the following parameters:

* -rights,-advanced-rights,-rights-raw
* -apply-to



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security descriptor
containing the system access control list entry whose fields you want to
modify.

``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor that contains the system
access control list entry that you want to modify.

``-access-type {failure|success}`` - Success or Failure::
Specifies whether the system access control list entry that you want to
modify is a ``_failure_`` or ``_success_`` access audit type.

``-account <name or sid>`` - Account Name or SID::
Specifies the account on which to apply the system access control list
entry. You can specify the account by using a user name or SID. You can
use any of the following formats when specifying the value for this
parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN
```

```
+
NOTE: If you specify any of the three user name formats for the value of
`-account` , keep in mind that the value for the user name is case
insensitive.
+


{ ``[-rights {no-access|full-control|modify|read-and-
execute|read|write}]`` - Access Rights::
Specifies the rights that you want to get audited for the account
specified in the `-account` parameter. The `-rights` parameter is mutually
exclusive with the `-advanced-rights` and `-rights-raw` parameter. If you
specify the `-rights` parameter, you can only specify one value.
+
You can specify one of the following rights values:
+
  * no-access
  * full-control
  * modify
  * read-and-execute
  * read
  * write



| ``[-rights-raw <Hex Integer>]`` - Raw Access Rights [.advaced-
priv]#(privilege: advanced)#::
Specifies the raw rights that you want to get audited for the account
specified in the `-account` parameter. The `-rights-raw` parameter is
mutually exclusive with the `-advanced-rights` and `-rights` parameter.
Specify the value as a hexadecimal integer, for example: ``_0xA10F_`` or
``_0xb3ff_`` etc.

| ``[-advanced-rights <Advanced access right>,...]`` - Advanced Access
Rights }::
Specifies the advanced rights that you want to get audited for the account
specified in the `-account` parameter. The `-advanced-rights` parameter is
mutually exclusive with the `-rights` and `-rights-raw` parameter. You can
specify more than one advanced-rights value by using a comma-delimited
list.
+
You can specify one or more of the following advanced rights:
+
  * read-data
  * write-data
```

```
  * append-data
  * read-ea
  * write-ea
  * execute-file
  * delete-child
  * read-attr
  * write-attr
  * delete
  * read-perm
  * write-perm
  * write-owner
  * full-control



``[-apply-to {this-folder|sub-folders|files}]`` - Apply SACL To::
Specifies where to apply the system access control list entry. You can
specify more than one value by using a comma-delimited list.
+
You can specify one or more of the following values:
+
  * this-folder
  * sub-folder
  * files


+
NOTE: Select one of the following combinations of values for the `-apply-
to` parameter for Storage-Level Access Guard (SLAG):
+
  * this-folder, sub-folder, files
  * this-folder, sub-folder
  * files


+
If you specify an invalid `-apply-to` value, this security descriptor is
removed from the associated Storage-Level Access Guard (SLAG) `security
file-directory policy task` .
+


== Examples

The following example modifies the rights and -apply-to fields in the SACL
entry.
```

```
cluster1::> vserver security file-directory ntfs sacl modify -ntfs-sd sd2 -access-type success -account
BUILTIN\Administrators -vserver vs1 -rights modify -apply-to this-folder,sub-folders
cluster1::> vserver security file-directory ntfs sacl show -vserver vs1 -ntfs-sd sd2 -account
BUILTIN\Administrators -instance
Vserver: vs1
Security Descriptor Name: sd2
Access type for Specified Access Rights: success
Account Name or SID: BUILTIN\Administrators
Access Rights: modify
Advanced Access Rights: -
Apply To: this-folder, sub-folders
Access Rights: modify
```

```
[[ID52df87a609aef62b693bc982fe9c286d]]
= vserver security file-directory ntfs sacl remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a SACL entry from NTFS security descriptor

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs sacl remove` command removes a
system access control list entry from a security descriptor.

You can unambiguously define which SACL entry to remove by specifying the
following four parameters in the command:

* Vserver associated with the security descriptor that contains the SACL
entry
* Name of the security descriptor that contains the SACL entry
* Whether the SACL is a success or failure type of SACL entry
* The account name or SID to which the SACL is applied



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security descriptor
from which you want to remove the system access control list entry.
```

```
``-ntfs-sd <ntfs sd name>`` - NTFS Security Descriptor Name::
Specifies the name of the security descriptor that contains the system
access control list entry that you want to remove.

``-access-type {failure|success}`` - Success or Failure::
Specifies whether the system access control list entry that you want to
remove is a ``_failure_`` or ``_success_`` access audit type.

``-account <name or sid>`` - Account Name or SID::
Specifies the account name or SID associated with the system access
control list entry that you want to remove.

== Examples

The following example removes a SACL entry named “sd2” on Vserver vs1 with
an access type of “success” associated with the "BUILTIN\Administrators"
account.
```

cluster1::> vserver security file-directory ntfs sacl remove -ntfs-sd sd2 -access-type success -account BUILTIN\Administrators -vserver vs1

```
[[ID4bb35e4f06d6bbf3e36a9b59b3951a76]]
= vserver security file-directory ntfs sacl show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NTFS security descriptor SACL entries

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory ntfs sacl show` command displays
information about all the system access control list entries in the
Vserver. The command output depends on the parameter or parameters
specified with the command. If you do not specify any parameters, the
command displays the following information about all SACL entries:

* Vserver name
* Security descriptor
* List of SACL entries
```

```
You can specify the `-fields` parameter to specify which fields of
information to display about SACL entries.

You can specify the `-instance` parameter to display all information about
SACL entries in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
system access control list entries associated with the specified Vserver.

``[-ntfs-sd <ntfs sd name>]`` - NTFS Security Descriptor Name::
If you specify this parameter, the command displays information only about
the system access control list entries for the security descriptor that
you specify.

``[-access-type {failure|success}]`` - Success or Failure::
If you specify this parameter, the command displays information only about
the system access control list entries with the access type that you
specify.

``[-account <name or sid>]`` - Account Name or SID::
If you specify this parameter, the command displays information only about
the system access control list entries associated with the account name or
SID that you specify. You can use any of the following formats when
specifying the value for this parameter:
+
+
  * SID
  * Domain\user-name
  * user-name@Domain
  * user-name@FQDN



+
NOTE: If you specify any of the three user name formats for the value of
```

```
-account, keep in mind that the value for the user name is case
insensitive.
+


``[-rights {no-access|full-control|modify|read-and-execute|read|write}]``
- Access Rights::
If you specify this parameter, the command displays information only about
the system access control list entries with the user right that you
specify. The value for this parameter is mutually exclusive with any other
rights values. Only one value can be specified.
+
You can specify one of the following rights values:
+
  * no-access
  * full-control
  * modify
  * read-and-execute
  * read
  * write



``[-rights-raw <Hex Integer>]`` - Raw Access Rights [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the system access control list entries with the advanced user rights that
you specify. This value for this parameter is mutually exclusive with any
other rights values. Specify the value as a hexadecimal integer, for
example: ``_0xA10F_`` or ``_0xb3ff_`` etc.

``[-advanced-rights <Advanced access right>,...]`` - Advanced Access
Rights::
If you specify this parameter, the command displays information only about
the system access control list entries with the advanced user rights that
you specify. You can specify more than one value by using a comma-
delimited list.
+
You can specify one or more of the following advanced rights values:
+
  * read-data
  * write-data
  * append-data
  * read-ea
  * write-ea
  * execute-file
```

```
  * delete-child
  * read-attr
  * write-attr
  * delete
  * read-perm
  * write-perm
  * write-owner
  * full-control



``[-apply-to {this-folder|sub-folders|files}]`` - Apply SACL To::
If you specify this parameter, the command displays information only about
the system access control list entries with the -applied-to value or
values that you specify. You can specify more than one value by using a
comma-delimited list.
+
You can specify one or more of the following values:
+
  * this-folder
  * sub-folder
  * files



``[-readable-access-rights <TextNoCase>]`` - Access Rights::
If you specify this parameter, the command displays information only about
the system access control list entries with the readable access rights
that you specify.

== Examples

The following example shows a SACL entry.
```

cluster1::> vserver security file-directory sacl show
(vserver security file-directory ntfs sacl show)
Vserver: vs1
NTFS Security Descriptor Name: sd1
Account Name Access Access Apply To
Type Rights
--- --- --- ---
domain\user success full-control this-folder, sub-folders, files

```
[[ID7172d4cfc6cbc3ffec7ab4b862cff89a]]
= vserver security file-directory policy create
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a file security policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory policy create` command creates a
security policy for a Vserver. A policy acts as a container for various
tasks where each task is a single entry that can be applied to a
file/folder.

Creating a security policy is the third step in configuring and applying
security ACLs to a file or folder. You will later add tasks to the
security policy.

NOTE: You cannot modify a security policy. If you want to apply a policy
with the same settings to a different Vserver, you must create a new
policy with the same configuration and apply it to the desired Vserver.



The steps to creating and applying NTFS ACLs are the following:

* Create an NTFS security descriptor.
* Add DACLS and SACLS to the NTFS security descriptor.

NOTE: If you want to audit file and directory events, you must configure
auditing on the Vserver in addition to adding SACLs to the security
descriptor.


* Create a file/directory security policy.

This step associates the policy with a Vserver.
* Create policy tasks.

A policy task refers to a single operation to apply to a file (or folder)
or to a set of files (or folders). among other things, the task defines
which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver on which to create the security policy.

``-policy-name <Security policy name>`` - Policy Name::
Specifies the name of the security policy.

== Examples

The following example creates a security policy named “policy1” on Vserver
vs1.
```

cluster1::> vserver security file-directory policy create -policy-name policy1 -vserver vs1
cluster1::> vserver security file-directory policy show
Vserver Policy Name
---------- ----------
vs1 policy1

```
[[ID12f90dac4faabb8eb883f344154817aa]]
= vserver security file-directory policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a file security policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 The `vserver security file-directory policy delete` command deletes a
security policy from a Vserver.

NOTE: Deleting a policy fails if a job is currently running for the
specified policy.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the name of the Vserver associated with the security policy that
you want to delete.

``-policy-name <Security policy name>`` - Policy Name::
Specifies the name of the security policy you want to delete.

== Examples

The following example deletes a security policy named “policy1” from
Vserver vs1.
```

cluster1::> vserver security file-directory policy delete -policy-name policy1 -vserver vs1

```
[[ID24000b8905f6504c96d84f9869d4b3fb]]
= vserver security file-directory policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display file security policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory policy show` command displays
information about all security policies in the Vserver. The command output
depends on the parameter or parameters specified with the command.

You can specify the `-fields` parameter to specify which fields of
information to display about security policies.

You can specify the `-instance` parameter to display information for all
security policies in a list format.

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
security policies associated with the specified Vserver.

``[-policy-name <Security policy name>]`` - Policy Name::
If you specify this parameter, the command displays information only about
the security policy you specify.

== Examples

The following example displays information about the security policies on
the cluster.
```

```
cluster1::> vserver security file-directory policy show
Vserver Policy Name
------------ --------------
vs1 policy1
vs1 policy2
```

2 entries were displayed.

```
[[ID253d042a326fac093cc6b9f1a594bec2]]
= vserver security file-directory policy task add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a policy task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory policy task add` command adds a
single task entry to a security policy. A task refers to a single
operation that can be done by a security policy to a file/folder.

Before you create a security policy task, you must first create a security
policy and a security descriptor. You should also add DACL entries and
SACL entries (if desired) to the security descriptor before you create the
security policy task.

NOTE: You can add DACL and SACL entries to the security descriptor after
you have associated it to a security policy task.



Creating a policy task is the fourth step in configuring and applying ACLs
to a file or folder. When you create the policy task, you associate a
security descriptor to it. You also associate the task to a security
policy.

The steps to creating and applying NTFS ACLs are the following:

* Create an NTFS security descriptor.
* Add DACLS and SACLS to the NTFS security descriptor.

NOTE: If you want to audit file and directory events, you must configure
auditing on the Vserver in addition to adding SACLs to the Security
Descriptor.
```

```
* Create a file/directory security policy.

This step associates the policy with a Vserver.
* Create policy tasks.

A policy task refers to a single operation to apply to a file (or folder)
or to a set of files (or folders). among other things, the task defines
which security descriptor to apply to a path.

NOTE: Adding a policy task fails if a job is currently running for the
specified policy to which a task is being added.


* Apply a policy to the associated Vserver.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver associated with the security policy to which you
want to add a task.

``-policy-name <Security policy name>`` - Policy Name::
Specifies the name of the security policy into which you want to add the
task.

``-path <text>`` - Path::
Specifies the path of the file/folder on which to apply the security
descriptor associated with this task.

``[-index-num <integer>]`` - Position::
Specifies the index number of a task. Tasks are applied in order. A task
with a larger index value is applied after a task with a lower index
number. If you do not specify this optional parameter, new tasks are
applied to the end of the index list.
+
The range of supported values is 1 through 9999. If there is a gap between
the highest existing index number and the value entered for this
parameter, the task with this number is considered to be the last task in
the policy and is treated as having an index number of the previous
highest index plus one.
+
NOTE: If you specify an index number that is already assigned to an
existing task, index number will be auto arranged to highest index number
in the table.
+
```

```
``[-security-type {ntfs|nfsv4}]`` - Security Type of the File::
Specifies whether the security descriptor associated with this task is an
NTFS or a NFSv4 security descriptor type. If you do not specify a value
for this optional parameter, the default is “ntfs”.
+
NOTE: The nfsv4 security descriptor type is not supported in this release.
If you specify this optional parameter, you must enter ntfs for the
-security-type value.
+


``[-ntfs-mode {propagate|ignore|replace}]`` - Propagation Mode::
Specifies how to propagate security settings to child subfolders and
files. This setting determines how child files and/or folders contained
within a parent folder inherit access control and audit information from
the parent folder.
+
You can specify one of the three parameter values that correspond to three
types of propagation modes:
+
  * propagate - propagate inheritable permissions to all subfolders and
files
  * replace - replace existing permissions on all subfolders and files
with inheritable permissions
  * ignore - do not allow permissions on this file or folder to be
replaced



+
NOTE: The ntfs-mode value is ignored for Storage-Level Access Guard
(SLAG).
+


``[-ntfs-sd <ntfs sd name>,...]`` - NTFS Security Descriptor Name::
Specifies the list of security descriptor names to apply to the path
specified in the `-path` parameter.

``[-access-control {file-directory|slag}]`` - Access Control Level::
Specifies the access control of the task to be applied. Valid values are
``_file-directory_`` or ``_slag_`` . Use the value ``_slag_`` to apply the
specified security descriptors with the task for the volume or qtree.
Otherwise, the security descriptors are applied on files and directories
at the specified path. The value ``_slag_`` is not supported on FlexGroup
```

```
volumes. The default value is ``_file-directory_`` .

== Examples

The following example adds a security policy task entry to the policy
named “policy1” on Vserver vs1.
```

cluster1::> vserver security file-directory policy task add -vserver vs1 -policy-name policy1 -path / -access -control slag -security-type ntfs -ntfs-mode propagate -ntfs-sd sd -index-num 1
cluster1::> vserver security file-directory policy task add -vserver vs1 -policy-name policy2 -path /1 -security -type ntfs -ntfs-mode propagate -ntfs-sd sd1,sd2
cluster1::> vserver security file-directory policy task show
Vserver: vs1
Policy: policy1
Index File/Folder Access Security NTFS NTFS Security
Path Control Type Mode Descriptor Name
----- ---------- --------- -------- ---------- ---------
1 / slag ntfs propagate sd
Vserver: vs1
Policy: policy2
Index File/Folder Access Security NTFS NTFS Security
Path Control Type Mode Descriptor Name
----- ---------- --------- -------- ---------- ---------
1 /1 file-directory ntfs propagate sd1, sd2

```
[[ID3945d9687018a9ab94ada53df36980ec]]
= vserver security file-directory policy task modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify policy tasks

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory policy task modify` command modifies
a task entry in a security policy.

NOTE: Modifying a policy task fails if a job is currently running for the
specified policy in which a task is being modified.
```

```
You can unambiguously define which task to modify by specifying the
following three parameters in the modify command:

* Vserver associated with the task
* Name of the security policy that contains the task
* Name of the path to which the task is applied



You can modify the following parameters:

* -ntfs-mode
* -ntfs-sd
* -index-num



NOTE: The only security type supported in this ONTAP release is
``_"ntfs"_`` ; therefore, you cannot modify the `-security-type`
parameter.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver associated with the security policy that contains
the task you want to modify.

``-policy-name <Security policy name>`` - Policy Name::
Specifies the name of the security policy that contains the task you want
to modify.

``-path <text>`` - Path::
Specifies the path of the file/folder associated with the task that you
want to modify.

``[-index-num <integer>]`` - Position::
Specifies the index number of a task. Tasks are applied in order. A task
with a larger index value is applied after a task with a lower index
number. If you do not specify this optional parameter, new tasks are
applied to the end of the index list.
+
The range of supported values is 1 through 9999. If there is a gap between
the highest existing index number and the value entered for this
parameter, the task with this number is considered to be the last task in
the policy and is treated as having an index number of the previous
```

```
highest index plus one.
+
NOTE: If you specify an index number that is already assigned to an
existing task, the command fails when you attempt to create a duplicate
entry.
+


``[-security-type {ntfs|nfsv4}]`` - Security Type::
Specifies whether the security descriptor in the task that you want to
modify should be an NTFS security descriptor type or an NFSv4 security
descriptor type. Default value is ``_ntfs_`` .
+
NOTE: The nfsv4 security descriptor type is not supported in this release.
If you specify this optional parameter, you must enter ntfs for the
-security-type value.
+


``[-ntfs-mode {propagate|ignore|replace}]`` - NTFS Propagation Mode::
Specifies how to propagate security settings to child subfolders and
files. This setting determines how child files and/or folders contained
within a parent folder inherit access control and audit information from
the parent folder.
+
You can specify one of the three parameter values that correspond to three
types of propagation modes:
+
  * propagate - propagate inheritable permissions to all subfolders and
files
  * replace - replace existing permissions on all subfolders and files
with inheritable permissions
  * ignore - do not allow permissions on this file or folder to be
replaced



``[-ntfs-sd <ntfs sd name>,...]`` - NTFS Security Descriptor Name::
Specifies the list of security descriptor names to apply to the path
specified in the `-path` parameter.

== Examples

The following example modifies the ntfs mode, index, and ntfs-sd
parameters in the security policy task entry.
```

```
cluster1::> vserver security file-directory policy task modify -vserver vs1 -policy-name policy1 -path / -security
-type ntfs -ntfs-mode propagate -ntfs-sd sd -index-num 1
cluster1::> vserver security file-directory policy task modify -vserver vs1 -policy-name policy1 -path /1 -security
-type ntfs -ntfs-mode propagate -ntfs-sd sd1, sd2 -index-num 2
cluster1::> vserver security file-directory policy task show -vserver vs1 -policy-name policy1
Vserver: vs1
Policy: policy1
Index File/Folder Access Security NTFS NTFS Security
Path Control Type Mode Descriptor Name
----- -------- -------------- ------ ----- -----------------
1 / file-directory ntfs propagate sd
2 /1 file-directory ntfs propagate sd1, sd2
```

```
[[ID81d19a3d295e7b03515609d06040c707]]
= vserver security file-directory policy task remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a policy task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory policy task remove` command removes a
task entry from a security policy.

NOTE: Removing a policy task fails if a job is currently running for the
specified policy from which a task is being removed.



== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver associated with the security policy that contains
the task you want to remove.

``-policy-name <Security policy name>`` - Policy Name::
Specifies the name of the security policy that contains the task you want
to remove.

``-path <text>`` - Path::
Specifies the path of the file/folder associated with the task that you
want to remove.

== Examples

The following example removes a security policy task entry.
```

cluster1::> vserver security file-directory policy task remove -vserver vs1 -policy-name policy1 -path /

```
[[ID0eaefee1313964ca85f6e35472e6a096]]
= vserver security file-directory policy task show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display policy tasks

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security file-directory policy task show` command displays
information about all the task entries in the Vserver. The command output
depends on the parameter or parameters specified with the command. If you
do not specify any parameters, the command displays the following
information about all task entries:

* Vserver name
* Policy name
* Task entries



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the -fields <fieldname>, ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only tasks
associated with the specified Vserver.

``[-policy-name <Security policy name>]`` - Policy Name::
If you specify this parameter, the command displays information only about
tasks associated with the specified security policy.

``[-index-num <integer>]`` - Position::
If you specify this parameter, the command displays information only about
tasks assigned the index number that you specify.
```

```
``[-path <text>]`` - Path::
If you specify this parameter, the command displays information only about
tasks applied to the specified path.

``[-security-type {ntfs|nfsv4}]`` - Security Type::
If you specify this parameter, the command displays information only about
tasks associated with the specified security type.
+
NOTE: The nfsv4 security descriptor type is not supported in this release.
+


``[-ntfs-mode {propagate|ignore|replace}]`` - NTFS Propagation Mode::
If you specify this parameter, the command displays information only about
tasks configured with the NTFS propagation mode that you specify.

``[-ntfs-sd <ntfs sd name>,...]`` - NTFS Security Descriptor Name::
If you specify this parameter, the command displays information only about
the policy tasks associated with the NTFS security descriptor that you
specify.

``[-access-control {file-directory|slag}]`` - Access Control Level::
If you specify this parameter, the command displays information only about
tasks associated to the access control.

== Examples

The following example displays policy task entries for a policy named
“policy1” on Vserver vs1.
```

cluster1::> vserver security file-directory policy task show -vserver vs1 -policy-name policy1
Vserver: vs1
Policy: policy1
Index File/Folder Access Security NTFS NTFS Security
Path Control Type Mode Descriptor Name
----- ---------- --------- -------- ---------- ---------
1 /1 file-directory ntfs propagate sd1, sd2
2 /2 file-directory ntfs ignore -
2 entries were displayed.

```
[[IDf50336ae34995a74629048610b167188]]
= vserver security trace filter create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Create a security trace entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security trace filter create` command creates a security
trace filter entry. Prior to ONTAP 9.3, this feature was only supported
for CIFS. In ONTAP 9.3 and later, this feature is supported for both NFS
and CIFS.

NFS security trace filters are not supported for FlexGroup volumes, and
will only be applied to the FlexVol volumes within the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which the permission
trace is applied.

``-index <integer>`` - Filter Index::
This parameter specifies the index number you want to assign to the trace
filter. A maximum of 10 entries can be created. The allowed values for
this parameter are 1 through 10.

``[-protocols {cifs|nfs}]`` - Protocols::
This parameter specifies the protocols for which the permission trace is
created. If the `-protocols` parameter is not specified, the filter will
only apply to the CIFS protocol.

``[-client-ip <IP Address>]`` - Client IP Address to Match::
This parameter specifies the IP Address from which the user is accessing
the Vserver.

``[-path <TextNoCase>]`` - Path::
This parameter specifies the path to which permission tracing is applied.
The value can be the complete path, starting from the root of the share
(for a CIFS filter) or the root of the junction path (for an NFS filter)
that the client is accessing, or the value can be a part of the path that
the client is accessing. Use NFS style directory separators in the path
value.

{ ``[-windows-name <TextNoCase>]`` - Windows User Name::
This parameter specifies the Windows user name to trace. You can use any
of the following formats when specifying the value for this parameter:
```

```
+
  *  user_name
  *  domain\user_name



| ``[-unix-name <TextNoCase>]`` - UNIX User Name or User ID }::
This parameter specifies the UNIX user name to trace. It accepts UNIX user
ID only for NFS filters.

``[-trace-allow {yes|no}]`` - Trace Allow Events::
Security tracing can trace deny events and allow events. Deny event
tracing is always ON by default. Allow events can optionally be traced. If
set to yes, this option allows tracing of allow events. If set to no,
allow events are not traced.

``[-enabled {enabled|disabled}]`` - Filter Enabled::
This parameter specifies whether to enable or disable the filter. Filters
are enabled by default.

``[-time-enabled <integer>]`` - Minutes Filter is Enabled::
This parameter specifies a timeout for this filter, after which it is
deleted.

== Examples

The following example creates a security trace filter.
```

cluster1::> vserver security trace filter create -vserver vs0 -index 1 -time-enabled 120 -client-ip 10.72.205.207

```
The following examples create filters that include the -path option, these
filters are deleted when the time specified in the time enabled field
elapses.The default value for the time-enabled option is 60 min. If the
client is accessing a file with the path
\\server\sharename\dir1\dir2\dir3\file.txt, for a filter applicable to
CIFS, a complete path starting from the root of the share or a partial
path can be given as shown:
```

cluster1::> vserver security trace filter create -vserver vs0 -index 1 -path /dir1/dir2/dir3/file.txt

```

```

cluster1::> vserver security trace filter create -vserver vs0 -index 1 -path dir3/file.txt

```
Similarly, while creating a filter for NFS, if -path option is specified
and the client is accessing a file with path
/junction_path1/junction_path2/dir1/file.txt, a complete path starting
from the last junction path or a partial path can be given as shown:
```

cluster1::> vserver security trace filter create -vserver vs0 -index 1 -protocols nfs -path dir1/file.txt

```

```

cluster1::> vserver security trace filter create -vserver vs0 -index 1 -protocols nfs -path file.txt

```
The following example creates a filter that is applicable to both CIFS and
NFS.
```

cluster1::> vserver security trace filter create -vserver vs0 -index 1 -protocols cifs,nfs -unix-user root

```
[[IDe409245779433da0043c04702c3c9c2f]]
= vserver security trace filter delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a security trace entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security trace filter delete` command deletes a security
trace filter entry. Prior to ONTAP 9.3, this feature was only supported
for CIFS. In ONTAP 9.3 and later, this feature is supported for both NFS
and CIFS.

NFS security trace filters are not supported for FlexGroup volumes, and
will only be applied to the FlexVol volumes within the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which the tracing
filter entry that you want to delete is applied.

``-index <integer>`` - Filter Index::
This parameter specifies the index number for the filter that you want to
delete. You can display a list of the filter index numbers by using the
xref:{relative_path}vserver-security-trace-filter-show.html[vserver
security trace filter show] command.

== Examples

The following example deletes a security trace filter.
```

cluster1::> vserver security trace filter delete -vserver vs0 -index 1

```
== Related Links

* xref:{relative_path}vserver-security-trace-filter-show.html[vserver
security trace filter show]
```

```
[[ID552a70b2a49c790f85d0da78b34a9051]]
= vserver security trace filter modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a security trace entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security trace filter modify` command modifies a security
trace filter entry. Prior to ONTAP 9.3, this feature was only supported
for CIFS. In ONTAP 9.3 and later, this feature is supported for both NFS
and CIFS.

NFS security trace filters are not supported for FlexGroup volumes, and
will only be applied to the FlexVol volumes within the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which the permission
trace is applied.

``-index <integer>`` - Filter Index::
This parameter specifies the index number for the filter. A maximum of 10
entries can be created. The allowed values for this parameter are 1
through 10.

``[-protocols {cifs|nfs}]`` - Protocols::
This parameter specifies the protocols for which the permission trace is
created.

``[-client-ip <IP Address>]`` - Client IP Address to Match::
This parameter specifies the IP Address from which the user is accessing
the Vserver.

``[-path <TextNoCase>]`` - Path::
```

```
This parameter specifies the path to which permission tracing is applied.
The value can be the complete path, starting from the root of the share
(for a CIFS filter) or the root of the junction path (for an NFS filter)
that the client is accessing, or the value can be a part of the path that
the client is accessing. Use NFS style directory separators in the path
value.

{ ``[-windows-name <TextNoCase>]`` - Windows User Name::
This parameter specifies the Windows user name to trace. You can use any
of the following formats when specifying the value for this parameter:
+
  *  user_name
  *  domain\user_name



| ``[-unix-name <TextNoCase>]`` - UNIX User Name or User ID }::
This parameter specifies the UNIX user name to trace. It accepts UNIX user
ID only for NFS filters.

``[-trace-allow {yes|no}]`` - Trace Allow Events::
Security tracing can trace deny events and allow events. Deny event
tracing is always ON by default. Allow events can optionally be traced. If
set to yes, this option allows tracing of allow events. If set to no,
allow events are not traced.

``[-enabled {enabled|disabled}]`` - Filter Enabled::
This parameter specifies whether to enable or disable the filter. Filters
are enabled by default.

``[-time-enabled <integer>]`` - Minutes Filter is Enabled::
This parameter specifies a timeout for this filter, after which it is
deleted.

== Examples

The following example modifies a security trace filter.
```

cluster1::> vserver security trace filter modify -vserver vs0 -index 1 -time-enabled 120 -client-ip 10.72.205.207

```
The following examples modify filters that include the -path option. If
the client is accessing a file with the path
\\server\sharename\dir1\dir2\dir3\file.txt, for a filter applicable to
CIFS, a complete path starting from the root of the share or a partial
path can be given as shown:
```

cluster1::> vserver security trace filter modify -vserver vs0 -index 1 -path /dir1/dir2/dir3/file.txt

```
```

cluster1::> vserver security trace filter modify -vserver vs0 -index 1 -path dir3/file.txt

```
Similarly, for filters applicable to NFS, if -path option is specified and
the client is accessing a file with path
/junction_path1/junction_path2/dir1/file.txt, a complete path starting
from the last junction path or a partial path can be given as shown:
```

cluster1::> vserver security trace filter modify -vserver vs0 -index 1 -protocols nfs -path dir1/file.txt

```
```

cluster1::> vserver security trace filter modify -vserver vs0 -index 1 -protocols nfs -path file.txt

```
The following example modifies a filter that is applicable to both CIFS
and NFS.
```

cluster1::> vserver security trace filter modify -vserver vs0 -index 1 -protocols cifs,nfs -unix-user root -path file.txt

```
[[IDb018679187c2b3b6c3a8753340c4ad38]]
= vserver security trace filter show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a security trace entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security trace filter show` command displays information
about security trace filter entries. Prior to ONTAP 9.3, this feature was
only supported for CIFS. In ONTAP 9.3 and later, this feature is supported
for both NFS and CIFS.

NFS security trace filters are not supported for FlexGroup volumes, and
```

```
will only be applied to the FlexVol volumes within the specified Vserver.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays permission tracing
information only for filters applied to the specified Vserver.

``[-index <integer>]`` - Filter Index::
If you specify this parameter, the command displays permission tracing
information only for filters with the specified filter index number.

``[-protocols {cifs|nfs}]`` - Protocols::
If you specify this parameter, the command displays permission tracing
information only for filters applied to the specified protocols.

``[-client-ip <IP Address>]`` - Client IP Address to Match::
If you specify this parameter, the command displays permission tracing
information only for filters applied to the specified client IP address.

``[-path <TextNoCase>]`` - Path::
If you specify this parameter, the command displays permission tracing
information only for filters applied to the specified path.

``[-windows-name <TextNoCase>]`` - Windows User Name::
If you specify this parameter, the command displays permission tracing
information only for filters applied to the specified Windows user name.

``[-unix-name <TextNoCase>]`` - UNIX User Name or User ID::
If you specify this parameter, the command displays permission tracing
information only for filters applied to the specified UNIX user name or
user ID(for NFS specific filters).

``[-trace-allow {yes|no}]`` - Trace Allow Events::
If you specify this parameter, the command displays information only about
events that either trace or do not trace allow events, depending on the
value provided.
```

```
``[-enabled {enabled|disabled}]`` - Filter Enabled::
If you specify this parameter, the command displays information only about
filters that either are enabled or disabled, depending on the value
provided.

``[-time-enabled <integer>]`` - Minutes Filter is Enabled::
If you specify this parameter, the command displays information about the
time durations configured for filters during creation.

== Examples

The following example displays security trace filters for Vserver
``_vserver1_`` .
```

cluster1::> vserver security trace filter show
Vserver Index Client-IP Path Trace-Allow Windows-Name Protocol
-------- ----- ------------ ---------- ----------- ------------- --------
vserver1 1 - - no domain\user cifs
vserver1 2 192.168.2.3 - yes - cifs
vserver1 3 - /dir1/dir2/file no domain\ cifs
administrator
vserver1 4 - file yes - nfs
4 entries were displayed.

```
[[IDb6b79d090831911c3323ef246841c5ac]]
= vserver security trace trace-result delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete security trace results

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 Delete the specified security tracing event record.

== Parameters

``-node {<nodename>|local}`` - Node::
 This parameter specifies the cluster node on which the permission tracing
event that you want to delete occurred.

``-vserver <vserver name>`` - Vserver::
 This parameter specifies the Vserver on which the permission tracing
event that you want to delete occurred.

``-seqnum <integer>`` - Sequence Number::
This parameter specifies the sequence number of the log entry to be
deleted.

== Examples

 The following example deletes the security trace result record for the
Vserver ``_vserver_1_`` on node ``_Node_1_`` whose sequence number is
``_999_`` .
```

cluster1::> vserver security trace trace-result delete -vserver vserver_1 -node Node_1 -seqnum 999

```
[[ID486b52f17a9a912f30e6a88f081ebafd]]
= vserver security trace trace-result show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display security trace results

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver security trace trace-result show` command displays the list
of security trace event records stored on the cluster. These records are
generated in response to security trace filters that are created using the
xref:{relative_path}vserver-security-trace-filter-create.html[vserver
security trace filter create] command. The command output depends on the
parameter or parameters specified with the command. If you do not specify
any parameters, the command displays the following information about all
the security trace events generated since the filter was enabled:

* Vserver name
* Cluster node name
* Security trace filter index number
* User name
* Security style
* Path
* Reason

 You can specify additional parameters to display only information that
match those parameters. For example, to display information about events
that occurred for the user "guest", run the command with `-user-name`
parameter set to ``_guest_`` .

== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify this parameter, the command only displays the fields that
you specify.

| ``[-instance  ]`` }::
If you specify this parameter, the command displays detailed information
about all security trace events.

``[-node {<nodename>|local}]`` - Node::
If you specify this parameter, the command displays information only about
security trace events on the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
```

```
security trace events on the specified Vserver.

``[-seqnum <integer>]`` - Sequence Number::
If you specify this parameter, the command displays information only about
the security trace events with this sequence number.

``[-keytime <Date>]`` - Time::
If you specify this parameter, the command displays information only about
security trace events that occurred at the specified time.

``[-index <integer>]`` - Index of the Filter::
If you specify this parameter, the command displays information only about
security trace events that occurred as a result of the filter
corresponding to the specified filter index number.

``[-client-ip <IP Address>]`` - Client IP Address::
If you specify this parameter, the command displays information only about
security trace events that occurred as a result of file access from the
specified client IP address.

``[-path <TextNoCase>]`` - Path of the File Being Accessed::
If you specify this parameter, the command displays information only about
the security trace events that occurred as a result of file accesses to
the specified path.

``[-win-user <TextNoCase>]`` - Windows User Name::
If you specify this parameter, the command displays information only about
the security trace events that occurred as a result of file access by the
specified Windows user.

``[-security-style <security style>]`` - Effective Security Style On
File::
If you specify this parameter, the command displays information only about
the security trace events that occurred on file systems with the specified
security style. The allowed values for security style are the following:
+
  * SECURITY_NONE - Security not Set
  * SECURITY_UNIX_MODEBITS - UNIX and UNIX permissions
  * SECURITY_UNIX_ACL - UNIX and NFSv4 ACL
  * SECURITY_UNIX_SD - UNIX and NT ACL
  * SECURITY_MIXED_MODEBITS - MIXED and UNIX permissions
  * SECURITY_MIXED_ACL - MIXED and NFSv4 ACL
  * SECURITY_MIXED_SD - MIXED and NT ACL
  * SECURITY_NTFS_MODEBITS - NTFS and UNIX permissions
  * SECURITY_NTFS_ACL - NTFS and NT ACL
  * SECURITY_NTFS_SD - NTFS and NT ACL
```

```
  * SECURITY_UNIX - UNIX
  * SECURITY_MIXED - MIXED
  * SECURITY_NTFS - NTFS
  * SECURITY_MODEBITS - UNIX permissions
  * SECURITY_ACL - ACL
  * SECURITY_SD - SD



``[-result <TextNoCase>]`` - Result of Security Checks::
If you specify this parameter, the command displays information about the
security trace events that have the specified result. Access to a file or
a directory can be 'allowed' or 'denied'. Output from this command
displays the result as a combination of the reason for allowing or denying
access, the location where access is either allowed or denied, and the
access right for which the file operation is allowed or denied.
+
The following are the reasons why an access can be allowed:
+
+
  * Access is allowed because the operation is trusted and no security is
configured
  * Access is allowed because the user has UNIX root privileges
  * Access is allowed because the user has UNIX owner privileges
  * Access is allowed because UNIX implicit permission grants requested
access
  * Access is allowed because the CIFS user is owner
  * Access is allowed because the user has take ownership privilege
  * Access is allowed because there is no CIFS ACL
  * Access is allowed because CIFS implicit permission grants requested
access
  * Access is allowed because the security descriptor is corrupted and the
user is a member of the Administrators group
  * Access is allowed because the ACL is corrupted and the user is a
member of the Administrators group
  * Access is allowed because the user has UNIX permissions
  * Access is allowed because explicit ACE grants requested access
  * Access is allowed because the user has audit privileges
  * Access is allowed because the user has superuser credentials
  * Access is allowed because inherited ACE grants requested access
  * Access is allowed because storage-level access guard (SLAG) grants
requested access
  * Access is allowed because no central access policies applied
  * Access is allowed because no central access policies could be applied
from the corrupt SACL
  * Access is allowed because matching central access policy could not be
```

```
located
  * Access is allowed because no central access rules apply to the object
  * Access is allowed because skipped one or more corrupt central access
rules
  * Access is allowed because all evaluated central access rules grant
access



+
+
The following are the reasons why an access can be denied:
+
+
  *  Access is denied by UNIX permissions
  *  Access is denied by an explicit ACE
  *  Access is denied. The requested permissions are not granted by the
ACE
  *  Access is denied. The security descriptor is corrupted
  *  Access is denied. The ACL is corrupted
  *  Access is denied. The sticky bit is set on the parent directory and
the user is not the owner of file or parent directory
  *  Access is denied. The owner can be changed only by root
  *  Access is denied. The UNIX permissions/uid/gid/NFSv4 ACL can be
changed only by owner or root
  *  Access is denied. The GID can be set by owner to a member of its
legal group list only if 'Owner can chown' is not set
  *  Access is denied. The file or the directory has readonly bit set
  *  Access is denied. There is no audit privilege
  *  Access is denied. Enforce DOS bits blocks the access
  *  Access is denied. Hidden attribute is set
  *  Access is denied by an inherited ACE
  *  Access is denied as the volume is readonly or directory is a snapshot
  *  Access is denied. System attribute is not set in the request
  *  Access is denied by the storage-level access guard (SLAG)
  *  Access is denied, file is infected
  *  Access is denied. Central access policy DB not ready
  *  Access is denied. Central access rule is corrupt
  *  Access is denied. Central access rule explicitly denied access
  *  Access is denied. Matching central access policy not found
  *  Access is denied because the user does not have UNIX root privileges
  *  Access is denied because the UNIX user could not be mapped to a valid
NT user
  *  Access is denied because the UNIX permissions/uid/gid/NFSv4 ACL
cannot be set in an NTFS qtree
```

```
+
+
The command or the location at which access was denied or allowed are as
follows:
+
+
  *  while traversing the directory.
  *  while truncating the file.
  *  while creating the directory.
  *  while creating the file.
  *  while checking parent's mode bits during delete.
  *  while deleting the child.
  *  while checking for child-delete access on the parent.
  *  while reading security descriptor.
  *  while accessing the link.
  *  while creating the directory.
  *  while creating or writing the file.
  *  while opening existing file or directory.
  *  while setting the attributes.
  *  while traversing the directory.
  *  while reading the file.
  *  while reading the directory.
  *  while deleting the target during rename.
  *  while deleting the child during rename.
  *  while writing data in the parent during rename.
  *  while adding a directory during rename.
  *  while adding a file during rename.
  *  while updating the target directory during rename.
  *  while setting attributes.
  *  while writing to the file.
  *  while extending the coral file.
  *  while creating the vdisk file.
  *  while checking for stale locks before open.
  *  while deleting a file or a directory.
  *  while truncating a hidden file.
  *  while truncating a file.
  *  while truncating a system file.
  *  while appending to a file or setting a file attribute.
  *  while opening a file or directory for delete.
  *  while checking for permission on parent directory during create.
  *  while appending to the file.
  *  while creating the device file.
  *  while reading the user's access rights on an object.


+
```

```
+
The access rights for which the file operation is allowed or denied are as
follows:
+
+
  *  Append.
  *  Delete.
  *  Delete Child.
  *  Execute.
  *  Generic All.
  *  Generic Execute.
  *  Generic Read.
  *  Generic Write.
  *  Maximum Allowed.
  *  Read.
  *  Read Attributes.
  *  Read Control.
  *  Read EA.
  *  System Security.
  *  Synchronize.
  *  Write.
  *  Write Attributes.
  *  Write DAC.
  *  Write EA.
  *  Write Owner.
  *  None.



``[-unix-user <TextNoCase>]`` - UNIX User Name::
If you specify this parameter, the command displays information only about
the security trace events that occurred as a result of file access by the
specified UNIX user.

``[-session-id <integer>]`` - CIFS Session ID::
If you specify this parameter, the command displays information only about
the security trace events that occurred as a result of file access by the
specified CIFS session ID.

``[-share-name <TextNoCase>]`` - Accessed CIFS Share Name::
If you specify this parameter, the command displays information only about
the security trace events that occurred as a result of file access by the
specified CIFS share name.

``[-protocol {cifs|nfs}]`` - Protocol::
If you specify this parameter, the command displays information only about
```

```
the security trace events that occurred for the specified protocol.

``[-volume-name <TextNoCase>]`` - Accessed Volume Name::
If you specify this parameter, the command displays information only about
the security trace events that occurred as a result of file access by the
specified volume name.

== Examples

The following example displays information about security trace records:
```

cluster1::> vserver security trace trace-result show
Vserver: vserver_1

Node Index Filter Details Reason
--- --- --- ---
cluster1-01 1 Security Style: MIXED Access is allowed because
and NT ACL CIFS implicit permission
grants requested access
while opening existing
file or directory.
Access is granted for:
"Read Attributes"
Protocol: cifs
Share: sh1
Path: /stk/bit
Win-User: cifs1\
administrator
Unix-User: root
Session-ID: 58455810

1 entries were displayed.

```
The following example displays information about security trace records
for path /stk/bit/set:
```

cluster1::> vserver security trace trace-result show -path /stk/bit/set

Vserver: vserver_1

Node Index Filter Details Reason
--- --- --- ---
cluster1-01 1 Security Style: MIXED Access is allowed because
and UNIX permissions the user has UNIX root
privileges while opening
existing file or
directory.
Access is granted for: "Read"
Protocol: cifs
Share: sh1

Path: /stk/bit/set
Win-User: cifs1\
administrator
UNIX-User: root
Session-ID: 75435293758455810
cluster1-01 1 Security Style: MIXED Access is denied. The
and NT ACL requested permissions
are not granted by the
ACE while checking for
child-delete access on
the parent. Access is not
granted for: "Delete Child"
Protocol: cifs
Share: sh1
Path: /stk/bit/set
Win-User: cifs1\
administrator
UNIX-User: root
Session-ID: 75435293758455324
cluster1-01 1 Security Style: MIXED Access is allowed because
and NT ACL the CIFS user is owner.
Access is denied by an
explicit ACE while
setting the attributes.
Access is not granted for:
"Read Attributes"
Protocol: cifs
Share: sh1
Path: /stk/bit/set
Win-User: cifs1\
administrator
UNIX-User: root
Session-ID: 75435293758455324
3 entries were displayed.

```
The following example displays information about security trace records
for the protocol nfs:
```

cluster1::> vserver security trace trace-result show -protocol nfs
Vserver: vserver_1

Node Index Filter Details Reason
--- --- --- ---
cluster1-01 2 Security Style: UNIX Access is allowed because the
permissions user has UNIX root privileges
while setting attributes.
Protocol: nfs
Volume: testvol_flex
Share: -
Path: /f1
Win-User: -
UNIX-User: root

Session-ID: -
cluster1-01 2 Security Style: UNIX Access is allowed because the
permissions user has UNIX root privileges
while writing to the file.
Access is granted for: "Write"
Protocol: nfs
Volume: testvol_flex
Share: -
Path: /f1
Win-User: -
UNIX-User: root
Session-ID: -
cluster1-01 3 Security Style: UNIX Access is denied by UNIX
permissions permissions while creating
the file. Access is not
granted for: "Synchronize",
"Read Control", "Read
Attributes", "Execute",
"Write"
Protocol: nfs
Volume: testvol_flex
Share: -
Path: /d1/file
Win-User: -
UNIX-User: 1029
Session-ID: -
3 entries were displayed.

```
== Related Links

* xref:{relative_path}vserver-security-trace-filter-create.html[vserver
security trace filter create]




:leveloffset: -1


= vserver services commands

:leveloffset: +1


[[ID1da027d5d9437dfd002e232d3876bb4b]]
= vserver services access-check authentication get-claim-name
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Get the Name of a Claim

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication get-claim-name` command
obtains the display name for a given claim.

== Parameters

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
The name of the Vserver.

``-claim-cn <text>`` - Claim CN [.advaced-priv]#(privilege: advanced)#::
The claim ID of the claim display name.

== Examples

This example gets the display name of a claim for the CIFS server created
on Vserver vs2
```

cluster1::vserver services access-check*> authentication get-dc-info -node vsim1 -vserver vs2 -claim-cn ad://ext/accountExpires:88d065c21536d9fe

Name of claim ad://ext/accountExpires:88d065c21536d9fe: accountExpires

```
[[ID249dbbbee539cb0f1d272192b59b6a4c]]
= vserver services access-check authentication get-dc-info
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Get Domain Controller Information

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication get-dc-info` command
obtains information about one of the Domain Controllers (DC) for the
domain of which the CIFS server is a member. The information fetched is
the Forest and Domain of which the DC is a member, the NetBIOS name of the
Domain, the NetBIOS Hostname of the DC, the CIFS Server site, the CIFS
Client site, GUID of the domain and flags. Flags describe the features and
roles of the DC.

== Parameters

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
The name of the Vserver.

== Examples

This example gets the information about a Domain Controller for CIFS
server created on Vserver vs2.
```

cluster1::vserver services access-check*> authentication get-dc-info -node vsim2-d1-01 -vserver vs2
DC Information:

Forest: cifs.lab.netapp.com
Domain: cifs.lab.netapp.com
NetBIOS Name: CIFSLAB
NetBIOS Hostname: A7-6
Server Site: cifs-dev-j4

Client Site:
GUID: 0366BE1F-FA08-4747-B5AC56097189C90E
Flags: 0x00000178
DS_LDAP_FLAG
DS_DS_FLAG
DS_KDC_FLAG
DS_TIMESERV_FLAG
DS_WRITABLE_FLAG
DS_PING_FLAGS

```
[[ID7c7fab930b502ba717b66f042965d1f7]]
= vserver services access-check authentication login-cifs
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Authenticate a CIFS user

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication login-cifs` command
authenticates a Windows user through the specified node using the
specified Vserver's configuration. Upon success, it displays the user's
Windows and UNIX credentials.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of a Vserver with a configured CIFS server.

``-user <text>`` - Windows Name [.advaced-priv]#(privilege: advanced)#::
The name of a user that is a member of the Windows or a trusted domain
that the CIFS server in the specified Vserver belongs to.

``[-clientIp <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
The IP address of the client as specified by the user

== Examples

This example authenticates the Windows user "administrator" through node
"node2" using the configuration of "vs1." Upon success, it displays the
Windows and UNIX credentials for user "administrator."
```

cluster1::vserver services access-check*> authentication login-cifs -vserver vs1 -user administrator -node node2

Enter the password:
Windows User: administrator Domain: EXAMPLE Privs: a7
Primary Grp: S-1-5-21-1407423728-2963865486-1834115207-513
Domain: S-1-5-21-1407423728-2963865486-1834115207 Rids: 500, 520, 513, 22226, 26625, 1842, 512, 519, 518, 8323, 1645, 1648, 1644, 1647
Domain: S-1-1 Rids: 0
Domain: S-1-5 Rids: 11, 2
Unix ID: 0, GID: 0
Flags: 1
Domain ID: 0
Other GIDs:
Authentication Succeeded.

```
[[ID75fc2fea165fbd2afb188a37c5e86400]]
= vserver services access-check authentication ontap-admin-ldap-fastbind
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Authenticate Data ONTAP admin CIFS user

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication ontap-admin-ldap-
fastbind` command authenticates the ONTAP administrator's user through the
specified node using LDAP fast bind mechanism. It uses the LDAP client
configuration of the specified Vserver.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of the Vserver with LDAP configuration.

``-user <text>`` - User Name [.advaced-priv]#(privilege: advanced)#::
The name of the user present in LDAP server.

== Examples

This example authenticates the ONTAP administrator user "administrator"
through node "node2" using the LDAP configuration of Vserver "vs1".
```

cluster1::vserver services access-check*> authentication ontap-admin-ldap-fastbind -vserver vs1 -user administrator -node node2

Enter the password:
Authentication Succeeded.

```
[[ID2aead430062909f0a20f9ff17a40086e]]
= vserver services access-check authentication ontap-admin-login-cifs
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Authenticate Data ONTAP admin CIFS user

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication ontap-admin-login-cifs`
command authenticates the ONTAP administrator's user through the specified
node using the specified Vserver's configuration. Upon success, it
displays the user's Windows credentials.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of a Vserver with a configured CIFS server.

``-user <text>`` - User Name [.advaced-priv]#(privilege: advanced)#::
The name of a user that is a member of the Windows or a trusted domain
that the CIFS server in the specified Vserver belongs to.

``[-clientIp <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
The IP address of the client as specified by the user

== Examples

This example authenticates the ONTAP administrator user "administrator"
through node "node2" using the configuration of Vserver "vs1." Upon
success, it displays the CIFS credentials for "administrator."
```

cluster1::vserver services access-check*> authentication ontap-admin-login-cifs -vserver vs1 -user administrator -node node2

Enter the password:
Windows User: administrator Domain: EXAMPLE Privs: a7
Primary Grp: S-1-5-21-1407423728-2963865486-1834115207-513
Domain: S-1-5-21-1407423728-2963865486-1834115207 Rids: 500, 520, 513, 22226, 26625, 1842, 512, 519, 518, 8323, 1645, 1648, 1644, 1647, 1003
Domain: S-1-1 Rids: 0
Domain: S-1-5 Rids: 11, 2
Authentication Succeeded.

```
[[IDbfda9d3d444f41a113a9cec848b6d092]]
= vserver services access-check authentication show-creds
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a user's credentials based on a UNIX UID or Windows SID or S3 User
Name

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication show-creds` command
returns the credentials for a Windows user using SID, a Windows user using
a Windows username, a UNIX user using UID, or a UNIX user using a UNIX
user name. This command is useful for retrieving information such as
account type, SIDs, UIDs, GIDs, privileges, and domain or group
membership.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The command displays information for the specified Vserver.

{ ``-uid <integer>`` - UID [.advaced-priv]#(privilege: advanced)#::
The UNIX user's UID.

| ``-sid <text>`` - SID [.advaced-priv]#(privilege: advanced)#::
The Windows user's SID.
```

```
| ``-unix-user-name <text>`` - Unix User Name [.advaced-priv]#(privilege:
advanced)#::
The UNIX username.

| ``-win-name <text>`` - Windows Name [.advaced-priv]#(privilege:
advanced)#::
The Windows username.

| ``-s3-user-name <text>`` - S3 User Name [.advaced-priv]#(privilege:
advanced)# }::
The S3 username.

``[-list-name {true|false}]`` - Display Translated Names [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information as
translated names.

``[-list-id {true|false}]`` - Display IDs [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information as IDs.

``[-clientIp <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
The IP address of the client as specified by the user

``[-skip-domain-group {true|false}]`` - Skip Domain Groups [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, Windows domain group membership will not
be fetched and only local group membership will be displayed, if any.

``[-show-partial-unix-creds {true|false}]`` - Display Partial UNIX
Credentials [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, partial UNIX credentials will be
displayed. This can be useful in cases where ONTAP is able to fetch the
UNIX credentials, but failed to fetch the Windows credentials.

== Examples

This example returns credential information for UNIX user with UID "0" on
node "node1" for Vserver "vs1."
```

cluster1::*> vserver services access-check authentication show-creds -node node1 -vserver vs1 -uid 0
(vserver services access-check authentication show-creds)
UNIX UID: root <> Windows User: CIFSQA\Administrator (User)
GID: root
Supplementary GIDs: <None>
Windows Membership:

CIFSQA\Schema Admins (Domain group)
CIFSQA\Enterprise Admins (Domain group)
CIFSQA\Domain Admins (Domain group)
CIFSQA\Domain Users (Domain group)
CIFSQA\Group Policy Creator Owners (Domain group)
BUILTIN\Administrators (Alias)
BUILTIN\Users (Alias)
User is also a member of Everyone, Authenticated Users, and Network Users
Privileges (0x2b7):
SeBackupPrivilege
SeRestorePrivilege
SeTakeOwnershipPrivilege

```
This example returns credential information for UNIX user with UID "0" on
node "node1" for Vserver "vs1" when list-name "false" and list-id "true."
```

cluster1::*> vserver services access-check authentication show-creds -node node1 -vserver vs1 -uid 0 -list -name false -list-id true
(vserver services access-check authentication show-creds)
UNIX UID: 0 <> Windows User: S-1-5-21-1407423728-2963865486-1834115207-500
GID: 0
Supplementary GIDs: <None>
Windows Membership:
S-1-5-21-1407423728-2963865486-1834115207-518
S-1-5-21-1407423728-2963865486-1834115207-519
S-1-5-21-1407423728-2963865486-1834115207-512
S-1-5-21-1407423728-2963865486-1834115207-513
S-1-5-21-1407423728-2963865486-1834115207-520
S-1-5-32-544
S-1-5-32-545
User is also a member of S-1-1-0, S-1-5-11, and S-1-5-2
Privileges (0x2b7):
SeBackupPrivilege
SeRestorePrivilege
SeTakeOwnershipPrivilege

```
This example returns credential information for UNIX user with UID "0" on
node "node1" for Vserver "vs1" when list-name "true" and list-id "true."
```

cluster1::*> vserver services access-check authentication show-creds -node node1 -vserver vs1 -uid 0 -list -name false -list-id true
(vserver services access-check authentication show-creds)
UNIX UID: 0 (root) <> Windows User: S-1-5-21-1407423728-2963865486-1834115207-500 (CIFSQA\Administrator (User))
GID: 0 (root)
Supplementary GIDs: <None>
Windows Membership:
S-1-5-21-1407423728-2963865486-1834115207-518 CIFSQA\Schema Admins (Domain group)
S-1-5-21-1407423728-2963865486-1834115207-519 CIFSQA\Enterprise Admins (Domain group)
S-1-5-21-1407423728-2963865486-1834115207-512 CIFSQA\Domain Admins (Domain group)
S-1-5-21-1407423728-2963865486-1834115207-513 CIFSQA\Domain Users (Domain group)

S-1-5-21-1407423728-2963865486-1834115207-520 CIFSQA\Group Policy Creator Owners (Domain group)
S-1-5-32-544 BUILTIN\Administrators (Alias)
S-1-5-32-545 BUILTIN\Users (Alias)
User is also a member of Everyone, Authenticated Users, and Network Users
Privileges (0x2b7):
SeBackupPrivilege
SeRestorePrivilege
SeTakeOwnershipPrivilege

```
This example returns credential information for UNIX user with UID "0" on
node "node1" for Vserver "vs1" when list-name "true" and list-id "false."
```

cluster1::*> vserver services access-check authentication show-creds -node node1 -vserver vs1 -uid 0 -list -name true -list-id false
(vserver services access-check authentication show-creds)
UNIX UID: root <> Windows User: CIFSQA\Administrator (User)
GID: root
Supplementary GIDs: <None>
Windows Membership:
CIFSQA\Schema Admins (Domain group)
CIFSQA\Enterprise Admins (Domain group)
CIFSQA\Domain Admins (Domain group)
CIFSQA\Domain Users (Domain group)
CIFSQA\Group Policy Creator Owners (Domain group)
BUILTIN\Administrators (Alias)
BUILTIN\Users (Alias)
User is also a member of Everyone, Authenticated Users, and Network Users
Privileges (0x2b7):
SeBackupPrivilege
SeRestorePrivilege
SeTakeOwnershipPrivilege

```
[[IDdcd23a98a651f0006007f02b7d88d583]]
= vserver services access-check authentication show-ontap-admin-unix-creds
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Data ONTAP admin Unix credentials based on username or user ID

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication show-ontap-admin-unix-
creds` uses the Vserver's ns-switch configuration to determine and display
ONTAP administrator's UNIX information.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of the Vserver.

{ ``-unix-user-name <text>`` - Unix User Name [.advaced-priv]#(privilege:
advanced)#::
The UNIX username.

| ``-uid <integer>`` - Unix User ID [.advaced-priv]#(privilege: advanced)#
}::
The UID of a UNIX user.

== Examples

This example shows ONTAP administrator's UNIX user's UID, GID, home
directory, and login shell for user "root" on Vserver "vs1" for node
"node2."
```

cluster1::vserver services access-check*> authentication show-ontap-admin-unix-creds -vserver vs1 -unix-user -name root -node node2
User ID: 0
Group ID: 1

Home Directory: /
Login Shell: /bin/csh

```
[[ID267487b088e4a46dd9061335516ee9df]]
= vserver services access-check authentication show-unix-ext-creds
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a user's UNIX extended credentials based on a UNIX UID

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication show-unix-ext-creds`
command returns the credentials for a UNIX user using UID. This command is
useful for retrieving information such as UIDs and GIDs.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The command displays information for the specified Vserver.

``-uid <integer>`` - UID [.advaced-priv]#(privilege: advanced)#::
The UNIX user's UID.

``[-ext-groups-limit <integer>]`` - Extended Groups Limit [.advaced-
priv]#(privilege: advanced)#::
The maximum number of GIDs to be fetched. The range is 32 to 1024. The
default value is 32.

``[-clientIp <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
The IP address of the client as specified by the user

== Examples

This example returns credential information for UNIX user with UID "100"
on node "node1" for Vserver "vs1."
```

cluster1::*> vserver services access-check authentication show-unix-ext-creds -node node1 -vserver vs1 -uid 100
(vserver services access-check authentication show-unix-ext-creds)
UNIX UID: 100 (user1)
UNIX GID: 200 (group1)
Additional GIDs (Count: 2): 200 201

```
[[ID32a761d3284c0d5c6319a77e00d9e7ee]]
= vserver services access-check authentication sid-to-uid
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Translate a Windows SID to a UNIX user ID

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication sid-to-uid` translates a
Windows SID to a UNIX UID.

== Parameters

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
The name of the Vserver.

``-sid <text>`` - Windows SID [.advaced-priv]#(privilege: advanced)#::
The SID of a Windows user.

``[-clientIp <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
The IP address of the client as specified by the user

== Examples

This example translates a Windows SID on node "node2" and returns the
corresponding UNIX user's UID.
```

cluster1::vserver services access-check*> sid-to-uid -vserver vs1 -sid S-1-5-21-1407423728-2963865486-1834115207-500 -node node2
UID: 0

```
[[IDd254b79e69cc09af8718152fd965f2b9]]
= vserver services access-check authentication sid-to-unix-name
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Translate a Windows SID to a UNIX User Name

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication sid-to-unix-name`
translates a Windows SID to a UNIX Name.

== Parameters

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
The name of the Vserver.

``-sid <text>`` - Windows SID [.advaced-priv]#(privilege: advanced)#::
The Windows SID which is to be translated to the corresponding UNIX name.

== Examples

This example translates a Windows SID on node "node2" and returns the
corresponding UNIX name.
```

cluster1::vserver services access-check*> sid-to-unix-name -node node2 -vserver vs1 -sid S-1-5-21-1407423728-2963865486-1834115207-500
SID Type: User
UNIX Name: test
Domain Name: TESTDOMAIN
Windows Name: test

```
[[ID4a53600518fe34d6b670fd69f7dff365]]
= vserver services access-check authentication translate
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Translate between Various Names and Their Identifiers

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication translate` command
translates SIDs, UIDs, and GIDs to names. If you enter a SID, the command
returns a Windows name; if you enter a Windows name, the command returns a
SID; if you enter a UNIX username, the command returns a UID; if you enter
a UID, the command returns a UNIX username; if you enter a GID, the
command returns a UNIX group name; if you enter a UNIX group-name, the
command returns a GID.

== Parameters

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
The name of the Vserver.

{ ``-uid <integer>`` - UNIX User ID [.advaced-priv]#(privilege:
advanced)#::
The UNIX user's UID.

| ``-gid <integer>`` - UNIX Group ID [.advaced-priv]#(privilege:
advanced)#::
The UNIX user's GID.

| ``-sid <text>`` - Windows SID [.advaced-priv]#(privilege: advanced)#::
The Windows user's SID.

| ``-unix-user-name <text>`` - UNIX User Name [.advaced-priv]#(privilege:
advanced)#::
The UNIX username.
```

```
| ``-unix-group-name <text>`` - UNIX Group Name [.advaced-
priv]#(privilege: advanced)#::
The UNIX group name.

| ``-win-name <text>`` - Windows Name [.advaced-priv]#(privilege:
advanced)# }::
The Windows name.

== Examples

This example translates the UNIX UID 0 to username "root" on node "node2"
for Vserver "vs1."
```

cluster1::vserver services access-check*> authentication translate -vserver vs1 -uid 0 -node node2
root

```
This example translates and the Windows username "administrator" to the
corresponding SID on node "node2" for Vserver "vs1."
```

cluster1::vserver services access-check*> authentication translate -vserver vs1 -win-name administrator -node
node2
S-1-5-21-1407423728-2963865486-1834115207-500

```
[[ID03af287fae90e6e7d5fad011434cebfc]]
= vserver services access-check authentication uid-to-sid
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Translate a UNIX user ID to a Windows SID

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check authentication uid-to-sid` translates a
UNIX UID to a Windows SID.

== Parameters

``[-node {<nodename>|local}]`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
The name of the Vserver.

``-uid <integer>`` - UNIX User ID [.advaced-priv]#(privilege: advanced)#::
The user ID of a UNIX user.

``[-clientIp <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
The IP address of the client as specified by the user

== Examples

This example translates a UNIX user's UID on node "node2" and returns the
corresponding SID.
```

cluster1::vserver services access-check*> uid-to-sid -vserver vs1 -uid 0 -node node2
SID: S-1-5-21-1407423728-2963865486-1834115207-500

```
[[ID36a355469a311f3f409d694fa2289f7a]]
= vserver services access-check dns forward-lookup
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Perform DNS forward lookup

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check dns forward-lookup` returns the IP
address of a hostname based on the lookup done on the DNS server specified
or the Vserver's DNS configuration.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
This specifies the name of the Vserver.

``-hostname <text>`` - Hostname [.advaced-priv]#(privilege: advanced)#::
This specifies the hostname to be looked up on the DNS server.

``[-lookup-type {ipv4|ipv6|all}]`` - Lookup Type (default: all) [.advaced-
priv]#(privilege: advanced)#::
This specifies the type of IP address to be looked up on the DNS server.
If you specify "all", it looks up both IPv4 and IPv6 addresses.

== Examples

The following example returns the IPv6 addresses of the hostname "example"
in Vserver "vs1" from the node "node2".
```

cluster1::vserver services access-check*> dns forward-lookup -vserver vs1 -node node2
-domains example.com -name-servers 10.72.46.234 -hostname example -lookup-type ipv6
6ffe::1
3ffe::1

```
[[ID4c3a1c19db413baca5c020d259887171]]
= vserver services access-check dns srv-lookup
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Perform DNS lookup for SRV records

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check dns srv-lookup` returns the IP
addresses of a host that is providing the specified service in the
network, based on the SRV record lookup done on the Vserver's DNS server.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
This specifies the name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
This specifies the name of the Vserver.

``[-name-servers <IP Address>,...]`` - (DEPRECATED)-Name Servers
[.advaced-priv]#(privilege: advanced)#::
This specifies the DNS servers in which the hostname lookup needs to be
done.

``-lookup-string <text>`` - Name to Lookup For [.advaced-priv]#(privilege:
advanced)#::
This specifies the complete string for which SRV record needs to be looked
up on the DNS server.

``[-lookup-type {ipv4|ipv6|all}]`` - Lookup Type (default: all) [.advaced-
priv]#(privilege: advanced)#::
This specifies the type of IP address to be looked up on the DNS server.
If you specify "all", it looks up both IPv4 and IPv6 addresses. The lookup
string must be in the form “service.protocol.domain”.

== Examples

The following example returns the IPv6 addresses of the host providing
```

```
 "http" service on "tcp" protocol in Vserver "vs1" from the node "node2".
```

cluster1::vserver services access-check*> dns srv-lookup -vserver vs1 -node node2
-lookup-string _http._tcp.nw7.na -lookup-type ipv6
9ffe::1
5ffe::1

```
[[ID040ae22de1ea6f2dc88ccb2c49a24d9f]]
= vserver services access-check name-mapping show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display or verify name mapping configuration

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check name-mapping show` command tests the
name mapping configuration for the specified Vserver on the specified
node. The command can perform name mapping for Kerberos to UNIX, Windows
to UNIX, UNIX to Windows, S3 to UNIX and S3 to Windows directions.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of the Vserver.

``-direction {krb-unix|win-unix|unix-win|s3-unix|s3-win}`` - Mapping
Direction [.advaced-priv]#(privilege: advanced)#::
The mapping direction.

``-name <text>`` - Name [.advaced-priv]#(privilege: advanced)#::
The username.

``[-clientIp <IP Address>]`` - Client IP Address [.advaced-
priv]#(privilege: advanced)#::
The IP address of the client as specified by the user

== Examples

This example shows a name mapping on Vserver "vs1" from UNIX username
"root," which is mapped to a Windows name "EXAMPLE\Administrator" on node
"node2."
```

cluster1::vserver services access-check*> name-mapping show -vserver vs1 -direction unix-win -name root -node node2
root maps to EXAMPLE\Administrator

```
This example shows a name mapping on Vserver "vs1" from Windows name
"EXAMPLE\Administrator" to a UNIX name "root."
```

cluster1::vserver services access-check*> name-mapping show -vserver vs1 -direction win-unix -name EXAMPLE\Administrator -node node2
EXAMPLE\Administrator maps to root

```
[[IDfd6eea7095ab24fc9f6e35800c6907d1]]
= vserver services access-check server-discovery reset
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Reset server discovery information for a Vserver

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check server-discovery reset` command deletes
all of the discovered server information for a given Vserver on the given
node. It returns a success message upon deleting. The next attempt to
access external servers will trigger the server discovery process to
acquire up-to-date server information.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of the Vserver on which you want to delete all of the discovered
server information.

== Examples

This example deletes all of the discovered server information for Vserver
"vs1" on the node "node2."
```

cluster1::vserver services access-check*> server-discovery reset -vserver vs1 -node node2
Discovery Reset succeeded for Vserver: vs1

```
[[ID27de9e0a0f7aa778a4a32d859fc588a0]]
= vserver services access-check server-discovery show-host
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display information about service host machines

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check server-discovery show-host` command
displays information about service host machines.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

== Examples

This example shows the host name and IP Address for five different service
host machines.
```

cluster1::vserver services access-check*> server-discovery show-host
Host Name: 172.19.3.11
Cifs Domain:
AD Domain:
IP Address: 172.19.3.11
Host Name: example-dc-1
Cifs Domain:
AD Domain:
IP Address: 172.17.152.40
Host Name: example-dc-2
Cifs Domain:
AD Domain:
IP Address: 172.17.152.41
Host Name: example-dc-3
Cifs Domain:
AD Domain:
IP Address: 172.17.152.42
Host Name: example-dc-4
Cifs Domain:

AD Domain:
IP Address: 172.17.152.43

```
[[IDa144c0dafa2d8c8abd22c29d96eca65f]]
= vserver services access-check server-discovery show-site
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display site membership

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services access-check server-discovery show-site` command
displays the site membership for a given Vserver on a given node.

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of the Vserver.

== Examples

This example shows the site membership for Vserver "vs1" from the
perspective of node "node2".
```

cluster1::vserver services access-check*> server-discovery show-site -node node2 -vserver vs1
california

```
[[IDccf6f4aa9c7a90ffd021c0d08bdb57b7]]
= vserver services access-check server-discovery test
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Verify server discovery

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The command `vserver services access-check server-discovery test`
completes the entire discovery including domain controller (DC), LDAP, and
NIS servers for the default domain for a given Vserver on a given node.
The Vserver must have CIFS configured for it to run successfully. If the
discovery is successful, the command returns a success message. The
discovered server information can be seen using the command `vserver cifs
domain discovered-servers` .

== Parameters

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
The name of the node on which the command is executed.

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
The name of the Vserver on which you want to test all of the discovered
server information.

== Examples

This example tests all of the discovered server information on Vserver
"vs1" on node "node2."
```

cluster1::*>vserver services access-check server-discovery test -vserver vs1 -node node2
Discovery Global succeeded for Vserver: vs1

```
[[ID4c2c5421357418053274eb9c10359fb5]]
= vserver services name-service cache group-membership delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The xref:{relative_path}vserver-services-name-service-cache-group-
membership-delete.html[vserver services name-service cache group-
membership delete] command removes the cached group membership entries of
the users for the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group membership
entries need to be deleted.

== Examples

The following example deletes all the cached group membership entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache group-membership delete-all -vserver vs0

```
== Related Links

* xref:{relative_path}vserver-services-name-service-cache-group-
membership-delete.html[vserver services name-service cache group-
membership delete]




[[ID0d94a9f70eef49a023037e009d727b8d]]
= vserver services name-service cache group-membership delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache group-membership delete` command
removes the cached group membership entries of the users.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group membership
entries need to be deleted.

``-user <text>`` - User Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the user name for which the cached group
membership entries need to be deleted.

``-group <integer>`` - Gid [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the primary group identifier or GID for
which the cached group membership entries need to be deleted.

== Examples

The following example deletes all the cached group membership entries for
Vserver vs0, user 'a' and group '1':
```

cluster1::> vserver services name-service cache group-membership delete -vserver vs0 -user a -group 1

```
[[IDabe86e83155378f65969104bed83ab7a]]
= vserver services name-service cache group-membership show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display group list

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache group-membership show` command
displays the cached group membership information of the users.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the cached group
membership details of the user.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group membership
entries need to be displayed.

``[-user <text>]`` - User Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached group
membership entries that have the specified user name.

``[-group <integer>]`` - Gid [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached group
membership entries of the users that have the specified primary group
identifier or GID.

``[-ngroups <integer>]`` - Number of Groups [.advaced-priv]#(privilege:
```

```
advanced)#::
Use this parameter to display information only about the cached group
membership entries of the users who belong to the specified number of
groups.

``[-groups <integer>,...]`` - Group List [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached group
membership entries of the users who belong to the specified group
identifiers or GIDs.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the group membership
entries that were cached at the specified time.

``[-is-partial {true|false}]`` - Is Partial Result [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the group membership
entries that have the specified value for partial result. The Value
``_true_`` displays only the cached entries that have partial result and
the value ``_false_`` displays only the cached entries that do not have
partial result.

== Examples

The following example displays the group membership details of the users
for all the vservers:
```

cluster1::> vserver services name-service cache group-membership show

```
The following example displays all the group membership details of the
users for Vserver vs0:
```

cluster1::> vserver services name-service cache group-membership show -vserver vs0

```
[[IDa8cd0650ad10128992ea7554b5c4ba45]]
= vserver services name-service cache group-membership settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify Group Membership Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache group-membership settings modify`
command modifies the group membership cache configuration of the specified
Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group membership
cache settings need to be modified.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
group membership database. The value ``_true_`` means the cache is enabled
and the value ``_false_`` means the cache is disabled. The default value
for this parameter is ``_true_`` .

``[-grplist-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Time to Live for Grplist [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes and seconds) for
which the group membership entries need to be cached. The default value is
1 hour.

== Examples

The following example modifies the group membership cache configuration
settings for Vserver vs0:
```

cluster1::> vserver services name-service cache group-membership settings modify -vserver vs0 -grplist-ttl 600

```
The following example disables the group membership cache for Vserver vs0:
```

cluster1::> vserver services name-service cache group-membership settings modify -vserver vs0 -is-enabled false

```
[[IDc2e050a224c9af832825d702cdffe918]]
= vserver services name-service cache group-membership settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Group Membership Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache group-membership settings show`
command displays information about the group membership cache
configuration for the users.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the group
membership cache configuration settings.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information about the group membership cache
configuration settings for the Vserver you specify.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the group membership
cache configuration settings that have the specified cache enabled
setting. The value ``_true_`` displays only the entries that have cache
```

```
enabled and the value ``_false_`` displays only the entries that have
cache disabled.

``[-grplist-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Time to Live for Grplist [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the group membership
cache configuration settings that have the specified Time to Live.

== Examples

The following example shows the group membership cache configuration
settings for all the Vservers:
```

cluster1::> vserver services name-service cache group-membership settings show

```
The following example shows the group membership cache configuration
settings for Vserver vs0:
```

cluster1::> vserver services name-service cache group-membership settings show -vserver vs0

```
[[ID99dc66d6e7af08365a16bfc519d9cbe9]]
= vserver services name-service cache hosts forward-lookup delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache hosts forward-lookup delete-all`
command removes all the cached host to IP lookup entries for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached forward
lookup entries need to be deleted.

== Examples

The following example deletes all the cached forward lookup entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache hosts forward-lookup delete-all -vserver vs0

```
[[IDe5ce24be48dfbe09cc694c622e5af959]]
= vserver services name-service cache hosts forward-lookup delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.
```

```
== Description

The `vserver services name-service cache hosts forward-lookup delete`
command removes a cached host to IP lookup entry.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached forward
lookup table entries need to be deleted.

``-host <text>`` - Hostname [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the hostname of the cached forward lookup
table entries that need to be deleted.

``-protocol {Any|ICMP|TCP|UDP}`` - Protocol [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the protocol of the cached forward lookup
table entries that need to be deleted.

``-sock-type {SOCK_ANY|SOCK_STREAM|SOCK_DGRAM|SOCK_RAW}`` - Sock Type
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the socket type of the cached forward lookup
table entries that need to be deleted.

``-flags
{FLAG_NONE|AI_PASSIVE|AI_CANONNAME|AI_NUMERICHOST|AI_NUMERICSERV}`` -
Flags [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the flag of the cached forward lookup table
entries that need to be deleted.

``-family {Any|Ipv4|Ipv6}`` - Family [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the family of the cached forward lookup
table entries that need to be deleted.

== Examples

The following example deletes the cached forward lookup entry for Vserver
vs0 and host "abc":
```

cluster1::> vserver services name-service cache hosts forward-lookup delete -vserver vs0 -host abc

```
[[ID4978ca0b32f6b0a7b9b94f56bb09ab08]]
= vserver services name-service cache hosts forward-lookup show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display host-byname struct

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache hosts forward-lookup show`
command displays the cached host to IP lookup entries.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the cached
forward lookup table entries.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached forward
lookup table entries need to be displayed.

``[-host <text>]`` - Hostname [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified hostname.

``[-protocol {Any|ICMP|TCP|UDP}]`` - Protocol [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified protocol.

``[-sock-type {SOCK_ANY|SOCK_STREAM|SOCK_DGRAM|SOCK_RAW}]`` - Sock Type
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified socket type.

``[-flags
```

```
{FLAG_NONE|AI_PASSIVE|AI_CANONNAME|AI_NUMERICHOST|AI_NUMERICSERV}]`` -
Flags [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified flags.

``[-family {Any|Ipv4|Ipv6}]`` - Family [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified family.

``[-canonname <text>]`` - Canonical Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified canonical name.

``[-ips <IP Address>,...]`` - IP Addresses [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified IPs.

``[-ip-protocol {Any|ICMP|TCP|UDP}]`` - Protocol [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified protocol of the resolved IP
address from forward lookup.

``[-ip-sock-type {SOCK_ANY|SOCK_STREAM|SOCK_DGRAM|SOCK_RAW}]`` - Sock Type
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified socket type of the resolved
IP address from forward lookup.

``[-ip-family {Any|Ipv4|Ipv6}]`` - Family [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified IP address family of the
resolved IP address from forward lookup.

``[-ip-addr-length <integer>,...]`` - Length [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified IP address length of the
resolved IP address from forward lookup.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
```

```
Use this parameter to display information only about the cached forward
lookup table entries that have the specified IP source of the resolved IP
address from forward lookup.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified time when the entry was
cached.

``[-ttl <integer>]`` - DNS TTL [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached forward
lookup table entries that have the specified Time To Live.

== Examples

The following example displays all the cached forward lookup entries:
```

cluster1::> vserver services name-service cache hosts forward-lookup show

```
The following example displays all the cached forward lookup entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache hosts forward-lookup show -vserver vs0

```
[[ID3c8ba77e12a3bd1a5f3cef1d941a63d2]]
= vserver services name-service cache hosts reverse-lookup delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache hosts reverse-lookup delete-all`
command removes all the cached IP to host lookup entries for a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver whose cached reverse lookup
entries need to be deleted.

== Examples

The following example deletes all the cached reverse lookup entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache hosts reverse-lookup delete-all -vserver vs0

```
[[ID85f76cdca9b7bd557cb2bf880a7308ce]]
= vserver services name-service cache hosts reverse-lookup delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache hosts reverse-lookup delete`
command removes a cached IP to host lookup entry.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached reverse
lookup table entries need to be deleted.

``-ip <IP Address>`` - IP Address [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the IP address of the cached reverse lookup
table entries that need to be deleted.

``-serv-flag <integer>`` - Service flags [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the service flag of the cached reverse
lookup table entries that need to be deleted.

== Examples

The following example deletes the cached reverse lookup entry for Vserver
vs0 and IP address 1.1.1.1:
```

cluster1::> vserver services name-service cache hosts reverse-lookup delete -vserver vs0 -ip 1.1.1.1

```
[[ID791ae5b75d20f9ff77d48f3ce15ca376]]
= vserver services name-service cache hosts reverse-lookup show
:hardbreaks:
:linkattrs:
```

```
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display ip-to-host struct

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache hosts reverse-lookup show`
command displays the cached IP to host lookup(reverse lookup) entries.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the cached
reverse lookup table entries.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached reverse
lookup table entries need to be displayed.

``[-ip <IP Address>]`` - IP Address [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified IP address.

``[-serv-flag <integer>]`` - Service flags [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified service flag.

``[-host <text>]`` - Hostname [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified hostname.

``[-service <text>]`` - Service Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached reverse
```

```
lookup table entries that have the specified service name.

``[-aliases <text>,...]`` - Host Aliases [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified aliases.

``[-addrtype <integer>]`` - Address Type [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified address type.

``[-addrlength <integer>]`` - Address Length [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified address length.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified create time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified source.

``[-ttl <integer>]`` - DNS TTL [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached reverse
lookup table entries that have the specified Time To Live.

== Examples

The following example displays all the cached reverse lookup entries:
```

cluster1::> vserver services name-service cache hosts reverse-lookup show

```
The following example displays the cached reverse lookup entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache hosts reverse-lookup show -vserver vs0

```
[[IDf59dc66794b4cf96320929097f701269]]
= vserver services name-service cache hosts settings modify
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify Hosts Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache hosts settings modify` command
modifies the hosts cache configuration of the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the hosts cache
settings need to be modified.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
hosts database. The value ``_true_`` means the cache is enabled and the
value ``_false_`` means the cache is disabled. The default value for this
parameter is ``_true_`` .

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
negative entries. Negative entries means the entries which are not present
in the hosts database and the lookup fails. The default value for this
parameter is ``_true_`` . Negative cache is disabled by default if the
parameter ``_is-enabled_`` is set to ``_false_`` .

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes and seconds) for
which the positive entries need to be cached. The positive entries means
the entries which are present in the hosts database and the lookup
succeeds. The default value is 24 hours.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time for which the negative entries need
```

```
to be cached. The default value is 1 minute.

``[-is-dns-ttl-enabled {true|false}]`` - Is TTL Taken from DNS [.advaced-
priv]#(privilege: advanced)#::
Specifies whether TTL is taken from DNS or host settings. If this
parameter is true, TTL is based on the DNS setting. If false, TTL is based
on the host setting. The default value is true.

== Examples

The following example modifies the hosts cache configuration settings for
Vserver vs0:
```

cluster1::> vserver services name-service cache hosts settings modify -vserver vs0 -ttl 600 -negative-ttl 300

```
The following example disables the cache for Vserver vs0:
```

cluster1::> vserver services name-service cache hosts settings modify -vserver vs0 -is-enabled false

```
[[ID616acfbe919e0a77109d4267a6a357cc]]
= vserver services name-service cache hosts settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Hosts Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache hosts settings show` command
displays information about the hosts cache configuration of the specified
Vserver.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the hosts cache
configuration settings.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information about the hosts cache
configuration settings for the Vserver you specify.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the hosts cache
configuration settings that have the specified cache enabled setting.
Value ``_true_`` displays only the entries that have cache enabled and
value ``_false_`` displays only the entries that have cache disabled.

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the hosts cache
configuration settings that have the specified negative cache enabled
setting. Value ``_true_`` displays only the entries that have negative
cache enabled and value ``_false_`` displays only the entries that have
negative cache disabled.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the hosts cache
configuration settings that have the specified Time to Live.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the hosts cache
configuration setings that have the specified negative Time to Live.

``[-is-dns-ttl-enabled {true|false}]`` - Is TTL Taken from DNS [.advaced-
priv]#(privilege: advanced)#::
Specifies whether TTL is based on the DNS or host settings.

== Examples

The following example shows the hosts cache configuration settings for
Vserver vs0:
```

cluster1::> vserver services name-service cache hosts settings show -vserver vs0

```
The following example shows the hosts cache configuration settings that
have cache disabled:
```

cluster1::> vserver services name-service cache hosts settings show -is-enabled false

```
[[ID8cce80fa14a3e2128ea9da2882e13efe]]
= vserver services name-service cache netgroups ip-to-netgroup delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups ip-to-netgroup delete-
all` command removes all the cached client IP to netgroup entries of the
specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached client IP
to netgroup entries need to be deleted.

== Examples

The following example deletes all the cached IP to netgroup entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache netgroups ip-to-netgroup delete-all -vserver vs0

```
[[ID1d3b1c9f893d9ffb1a108cee3b26d7b9]]
= vserver services name-service cache netgroups ip-to-netgroup delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete netgroup.byhost cache entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups ip-to-netgroup delete`
command removes the cached client IP to netgroup entries.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached client IP
to netgroup entries need to be deleted.

``-host <text>`` - Host field [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the IP address for which the cached IP to
netgroup entries need to be deleted.

``-netgrp <text>`` - Netgroup field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the netgroup for which the cached IP to
netgroup entries need to be deleted.

== Examples

The following example deletes all the cached IP to netgroup entries for
Vserver vs0, host 1.1.1.1 and netgrp 'abc':
```

cluster1::> vserver services name-service cache netgroups ip-to-netgroup delete -vserver vs0 -host 1.1.1.1 -netgrp abc

```
[[ID17cf2afb815ce622f289a1194c0d40a9]]
= vserver services name-service cache netgroups ip-to-netgroup show
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display netgroup.byhost cache entries

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups ip-to-netgroup show`
command displays the cached client IP to netgroup entries.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the cached client
IP to netgroup entries.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached client IP
to netgroup entries need to be displayed.

``[-host <text>]`` - Host field [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached IP to
netgroup entries that have the specified IP address.

``[-netgrp <text>]`` - Netgroup field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached IP to
netgroup entries that have the specified netgroup.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the IP to netgroup
entries that were cached at the specified time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
```

```
Use this parameter to display information only about the cached IP to
netgroup entries that have the specified look-up source.

== Examples

The following example displays all the cached IP to netgroup entries:
```

cluster1::> vserver services name-service cache netgroups ip-to-netgroup show

```
The following example shows all the cached IP to netgroup entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache netgroups ip-to-netgroup show -vserver vs0

```
[[ID9e49d31b7d89ee050eb92bdd1117f2e2]]
= vserver services name-service cache netgroups members delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups members delete-all`
command deletes all the cached netgroup member entries of the specified
Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached netgroup
members entries need to be deleted.

== Examples

The following example deletes all the cached netgroup members of Vserver
vs0:
```

cluster1::> vserver services name-service cache netgroups members delete-all -vserver vs0

```
[[ID18c159e25104e5b95d78a10468a399b3]]
= vserver services name-service cache netgroups members delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete netgroup cache entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups members delete` command
deletes the cached members of the netgroups.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached netgroup
members entries need to be deleted.

``-netgroup <text>`` - Netgroup [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the netgroup for which the cached netgroup
members entries need to be deleted.

== Examples

The following example deletes all the cached netgroup members entries for
Vserver vs0 and netgroup 'abc':
```

cluster1::> vserver services name-service cache netgroups members delete -vserver vs0 -netgroup abc

```
[[IDf1ebb55361bc4f2f19d0d72c126a9570]]
= vserver services name-service cache netgroups members show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display netgroup cache entries

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups members show` command
displays the cached members of the netgroups.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the cached
members of a netgroup.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the cached netgroup
members entries need to be displayed.

``[-netgroup <text>]`` - Netgroup [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached members
that belong to the specified netgroup.

``[-hosts <text>]`` - Hosts [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached netgroups
that have the specified host as a member.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the netgroup member
entries that were cached at the specified time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the cached netgroup
member entries that have the specified look-up source.

== Examples
```

```
The following example displays all the cached netgroup members entries:
```

cluster1::> vserver services name-service cache netgroups members show

```
The following example displays all the cached netgroup members entries for
Vserver vs0:
```

cluster1::> vserver services name-service cache netgroups members show -vserver vs0

```
[[ID8309d0746a4a24ef8559d7d6c718d603]]
= vserver services name-service cache netgroups settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify Netgroup Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups settings modify`
command modifies the netgroups cache configuration of the specified
Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the netgroups cache
settings need to be modified.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
netgroups database. The value ``_true_`` means the cache is enabled and
the value ``_false_`` means the cache is disabled. The default value for
this parameter is ``_true_`` .

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
```

```
negative entries. Negative entries means the entries which are not present
in the netgroups database and the look-up fails. The default value for
this parameter is ``_true_`` . Negative cache is disabled by default if
the parameter ``_is-enabled_`` is set to ``_false_`` .

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes and seconds) for
which the positive entries need to be cached. The positive entries means
the entries which are present in the netgroups database and the look-up
succeeds. The default value is 24 hours.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes and seconds) for
which the negative entries need to be cached. The default value is 30
minutes.

``[-ttl-members <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
TTL for netgroup members [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes and seconds) for
which the netgroup members need to be cached. The default value is 24
hours.

== Examples

The following example modifies the netgroups cache configuration settings
for Vserver vs0:
```

cluster1::> vserver services name-service cache netgroups settings modify -vserver vs0 -ttl 600 -negative-ttl 300

```
The following example disables the cache for Vserver vs0:
```

cluster1::> vserver services name-service cache netgroups settings modify -vserver vs0 -is-enabled false

```
[[ID1cc0874de529687fca6d53859a028cde]]
= vserver services name-service cache netgroups settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Netgroup Cache Configuration
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache netgroups settings show` command
displays information about the netgroups cache configuration of the
specified Vserver.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the netgroups
cache configuration settings.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information about the netgroups cache
configuration settings for the Vserver you specify.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the netgroups cache
configuration settings that have the specified cache enabled setting.
Value ``_true_`` displays only the entries that have cache enabled and
value ``_false_`` displays only the entries that have cache disabled.

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the netgroups cache
configuration settings that have the specified negative cache enabled
setting. Value ``_true_`` displays only the entries that have negative
cache enabled and value ``_false_`` displays only the entries that have
negative cache disabled.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the netgroups cache
configuration settings that have the specified Time to Live.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
```

```
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the netgroups cache
configuration setings that have the specified negative Time to Live.

``[-ttl-members <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
TTL for netgroup members [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the netgroups cache
configuration settings that have the specified Time to Live for netgroup
members.

== Examples

The following example shows the netgroups cache configuration settings for
Vserver vs0:
```

cluster1::> vserver services name-service cache netgroups settings show -vserver vs0

```
The following example shows the netgroups cache configuration settings
that have cache disabled:
```

cluster1::> vserver services name-service cache netgroups settings show -is-enabled false

```
[[ID95a233153934f6c2f14c6a0dbf9cd990]]
= vserver services name-service cache settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify nameservice global cache settings

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service cache settings modify` command modifies
the global name service cache configuration.

== Parameters

``[-eviction-time-interval
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Cache Eviction
Time Interval [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time interval after which a periodic
cache eviction will happen. The default value is 4 hours.

``[-is-remote-fetch-enabled {true|false}]`` - Is Remote Fetch Enabled
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify whether a node is allowed to fetch the data
from a remote node or not. If the value is set as ``_false_`` , the node
is not allowed to fetch the data from the remote node. If the value is set
as ``_true_`` , remote fetch is allowed.

== Examples

The following example modifies the global nameservice cache configuration:
```

cluster1::> vserver services name-service cache settings modify -eviction-time-interval 1h -is-remote-fetch -enabled true

```
[[IDa5e29dcaf02946c0dfaadc707dcbe91a]]
= vserver services name-service cache settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display nameservice global cache settings

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service cache settings show` command displays
information about the global name service cache configuration.

== Examples

The following example shows the global nameservice cache configuration:
```

cluster1::> vserver services name-service cache settings show

```
[[ID93bd6cb1dbae59bc6dc359dfd724920c]]
= vserver services name-service cache unix-group group-by-gid delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group group-by-gid delete-
all` command removes all the group entries that are cached by the group
identifier or GID.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group entries that
are cached by the group identifier or GID need to be deleted.

== Examples

The following example deletes all the cached group entries for Vserver
vs0:
```

cluster1::> vserver services name-service cache unix-group group-by-gid delete-all -vserver vs0

```
[[IDb246df58e00f815f8b3d769909fbcac8]]
= vserver services name-service cache unix-group group-by-gid delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group group-by-gid delete`
command removes the group entries that are cached by the group identifier
or GID. If group cache propagation is enabled, the corresponding group-by-
name cache entry will also be removed.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group entries that
are cached by the group identifier or GID need to be deleted.

``-gr-gid <integer>`` - gr_gid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the group identifier or GID for which the
cached entries need to be deleted.

== Examples

The following example deletes all the cached group entries for Vserver vs0
and the group identifier or GID 123:
```

cluster1::> vserver services name-service cache unix-group group-by-gid delete -vserver vs0 -gr-gid 123

```
[[IDacbd72ddf08044eca75a66f0c6285b52]]
= vserver services name-service cache unix-group group-by-gid show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display group struct

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group group-by-gid show`
command displays the group information cached by the group identifier or
GID.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the group entries
cached by the group identifier or GID.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group entries that
are cached by the group identifier or GID need to be displayed.

``[-gr-gid <integer>]`` - gr_gid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached group
entries that have the specified group identifier or GID.

``[-gr-name <text>]`` - gw_name field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached group
entries that have the specified group name.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the group entries
that were cached at the specified time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the group entries
```

```
cached by the group identifier or GID that have the specified lookup
source.

== Examples

The following example displays all the groups which are cached by the
group identifier or GID:
```

cluster1::> vserver services name-service cache unix-group group-by-id show

```
The following example displays all the group entries cached by the group
identifier or GID for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-group group-by-id show -vserver vs0

```
[[IDbfa7b0dfe94cc703442d1a650cfa8d22]]
= vserver services name-service cache unix-group group-by-name delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group group-by-name delete-
all` command removes all the group entries that are cached by the group
name for the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group entries that
are cached by group name need to be deleted.

== Examples

The following example deletes all the cached group entries for Vserver
vs0:
```

cluster1::> vserver services name-service cache unix-group group-by-name delete-all -vserver vs0

```
[[IDabfc62f92dc398183f4c2e9c59410f10]]
= vserver services name-service cache unix-group group-by-name delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group group-by-name delete`
command removes the group entries that are cached by group name. If group
cache propagation is enabled, the corresponding group-by-gid cache entry
will also be removed.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group entries that
are cached by group name need to be deleted.

``-gr-name <text>`` - gw_name field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the group name for which the cached entries
need to be deleted.

== Examples

The following example deletes all the cached group entries for Vserver vs0
and group name abc:
```

cluster1::> vserver services name-service cache unix-group group-by-name delete -vserver vs0 -gr-name abc

```
[[ID7e0a95b8db8fe5ca5db75d8aa70e3acb]]
= vserver services name-service cache unix-group group-by-name show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display group struct

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group group-by-name show`
command displays the group information cached by group name.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the group entries
cached by group name.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the group entries that
are cached by group name need to be displayed.

``[-gr-name <text>]`` - gw_name field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached group
entries that have the specified group name.

``[-gr-gid <integer>]`` - gr_gid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached group
entries that have the specified group identifier or GID.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the group entries
that were cached at the specified time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the group entries
cached by group name that have the specified lookup source.
```

```
== Examples

The following example displays all the groups which are cached by group
name:
```

cluster1::> vserver services name-service cache unix-group group-by-name show

```
The following example displays all the group entries cached by group name
for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-group group-by-name show -vserver vs0

```
[[ID824c39bc2d53bc354315707e56aa10c8]]
= vserver services name-service cache unix-group settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify UNIX Group Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group settings modify`
command modifies the groups cache configuration of the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the groups cache
settings need to be modified.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
groups database. The value ``_true_`` means the cache is enabled and the
value ``_false_`` means the cache is disabled. The default value for this
parameter is ``_true_`` .
```

```
``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
negative entries. Negative entries means the entries which are not present
in the groups database and the lookup fails. The default value for this
parameter is ``_true_`` . Negative cache is disabled by default if the
parameter ``_is-enabled_`` is set to ``_false_`` .

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time(in hours, minutes and seconds) for
which the positive entries need to be cached. The positive entries means
the entries which are present in the groups database and the lookup
succeeds. The default value is 24 hours.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time(in hours, minutes and seconds) for
which the negative entries need to be cached. The default value is 5
minutes.

``[-is-propagation-enabled {true|false}]`` - Is Propagation Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify whether the cached groups entries need to be
propagated to the group by the group identifier or GID cache. The default
value is ``_true_`` . Specify ``_false_`` to disable propagation.

== Examples

The following example modifies the groups cache configuration settings for
Vserver vs0:
```

cluster1::> vserver services name-service cache unix-group settings modify -vserver vs0 -ttl 600 -negative-ttl 300

```
The following example disables the cache for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-group settings modify -vserver vs0 -is-enabled false

```
[[ID985c5cfd0e7f56ec3171b82a86f547af]]
= vserver services name-service cache unix-group settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display UNIX Group Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-group settings show` command
displays information about the groups cache configuration of the specified
Vserver.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the groups cache
configuration settings.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information about the groups cache
configuration settings for the Vserver you specify.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the groups cache
configuration settings that have the specified cache enabled setting. The
value ``_true_`` displays only the entries that have cache enabled and the
value ``_false_`` displays only the entries that have cache disabled.

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the groups cache
configuration settings that have the specified negative cache enabled
setting. The value ``_true_`` displays only the entries that have negative
cache enabled and the value ``_false_`` displays only the entries that
have negative cache disabled.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the groups cache
```

```
configuration settings that have the specified Time to Live.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the groups cache
configuration setings that have the specified negative Time to Live.

``[-is-propagation-enabled {true|false}]`` - Is Propagation Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the groups cache
configuration settings that have the specified propagation enabled
setting. The value ``_true_`` displays only the entries that have the
propagation of cached entries to groups by the group identifier or GID
cache enabled and the value ``_false_`` displays only the entries that
have the propagation of cached entries to groups by the group identifier
or GID cache disabled.

== Examples

The following example shows the groups cache configuration settings for
all the Vservers:
```

cluster1::> vserver services name-service cache unix-group settings show

```
The following example shows the groups cache configuration settings for
Vserver vs0:
```

cluster1::> vserver services name-service cache unix-group settings show -vserver vs0

```
The following example shows the groups cache configuration settings that
have cache disabled:
```

cluster1::> vserver services name-service cache unix-group settings show -is-enabled false

```
[[IDa797d2c8a566f383071d864e5aa5273f]]
= vserver services name-service cache unix-user settings modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify UNIX users Cache Configuration
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user settings modify`
command modifies the users cache configuration of the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the users cache
settings need to be modified.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
users database.The value ``_true_`` means the cache is enabled and the
value ``_false_`` means the cache is disabled. The default value for this
parameter is ``_true_`` .

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify if the cache needs to be enabled for the
negative entries. Negative entries means the entries which are not present
in the users database and the look-up fails. The default value for this
parameter is ``_true_`` . Negative cache is disabled by default if the
parameter ``_is-enabled_`` is set to ``_false_`` .

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes and seconds) for
which the positive entries need to be cached. The positive entries means
the entries which are present in the users database and the look-up
succeeds. The default value is 24 hours.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the time (in hours, minutes and seconds) for
which the negative entries need to be cached. The default value is 5
minutes.

``[-is-propagation-enabled {true|false}]`` - Is Propagation Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify whether the cached users entries need to be
propagated to the users by id cache. The default value is ``_true_`` .
```

```
Specify ``_false_`` to disable propagation.

== Examples

The following example modifies the users cache configuration settings for
Vserver vs0:
```

cluster1::> vserver services name-service cache unix-user settings modify -vserver vs0 -ttl 600 -negative-ttl 300

```
The following example disables the cache for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-user settings modify -vserver vs0 -is-enabled false

```
[[IDbf0274f4c0660c3869fb193af70084bd]]
= vserver services name-service cache unix-user settings show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display UNIX users Cache Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user settings show` command
displays information about the users cache configuration of the specified
Vserver.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the users cache
configuration settings.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
```

```
advanced)#::
Use this parameter to display information about the users cache
configuration settings for the Vserver you specify.

``[-is-enabled {true|false}]`` - Is Cache Enabled? [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the users cache
configuration settings that have the specified cache enabled setting.
Value ``_true_`` displays only the entries that have cache enabled and
value ``_false_`` displays only the entries that have cache disabled.

``[-is-negative-cache-enabled {true|false}]`` - Is Negative Cache Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the users cache
configuration settings that have the specified negative cache enabled
setting. Value ``_true_`` displays only the entries that have negative
cache enabled and value ``_false_`` displays only the entries that have
negative cache disabled.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the users cache
configuration settings that have the specified Time to Live.

``[-negative-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Negative Time to Live [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the users cache
configuration setings that have the specified negative Time to Live.

``[-is-propagation-enabled {true|false}]`` - Is Propagation Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the users cache
configuration settings that have the specified propagation enabled
setting. Value ``_true_`` displays only the entries that have the
propagation of cached entries to users by id cache enabled and value
``_false_`` displays only the entries that have the propagation of cached
entries to users by id cache disabled.

== Examples

The following example shows the users cache configuration settings for all
the Vservers:
```

cluster1::> vserver services name-service cache unix-user settings show

```
The following example shows the users cache configuration settings for
Vserver vs0:
```

cluster1::> vserver services name-service cache unix-user settings show -vserver vs0

```
The following example shows the users cache configuration settings that
have cache disabled:
```

cluster1::> vserver services name-service cache unix-user settings show -is-enabled false

```
[[IDcaafcbb8ca2ed145059e285443775a97]]
= vserver services name-service cache unix-user user-by-id delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user user-by-id delete-all`
command removes all the user entries that are cached by the user
identifier or UID for the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the user entries that
are cached by the user identifier or UID need to be deleted.

== Examples

The following example deletes all the cached user entries for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-user user-by-id delete-all -vserver vs0

```
[[ID7f828d0aab72cd51d02985ada15f180e]]
= vserver services name-service cache unix-user user-by-id delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user user-by-id delete`
command removes the user entries that are cached by the user identifier or
UID. If user cache propagation is enabled, the corresponding user-by-name
cache entry will also be removed.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the user entries that
are cached by the user identifier or UID need to be deleted.

``-pw-uid <integer>`` - pw_uid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the user identifier or UID for which the
cached entries need to be deleted.

== Examples

The following example deletes all the user entries cached by the user
identifier or UID for Vserver vs0 and user identifier or UID 123:
```

cluster1::> vserver services name-service cache unix-user user-by-id delete -vserver vs0 -pw-uid 123

```
[[ID00106d666e651868141410e8520c20b7]]
= vserver services name-service cache unix-user user-by-id show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display password struct

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user user-by-id show`
command displays the user information cached by the user identifier or
UID.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the user entries
cached by the user identifier or UID.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the user entries that
are cached by the user identifier or UID need to be displayed.

``[-pw-uid <integer>]`` - pw_uid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the user identifier or UID for which the
cached entries need to be displayed.

``[-pw-name <text>]`` - pw_name field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached user
entries that have the specified user identifier or UID.

``[-pw-gid <integer>]`` - pw_gid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached user
entries that have the specified group identifier or GID.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the user entries that
```

```
were cached at the specified time.

``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the user entries
cached by the user identifier or UID that have the specified lookup
source.

== Examples

The following example displays all the users which are cached by the user
identifier or UID:
```

cluster1::> vserver services name-service cache unix-user user-by-id show

```
The following example displays all the users entries cached by the user
identifier or UID for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-user user-by-id show -vserver vs0

```
[[ID2f586fdd329d4f45c752ad45dd732fb9]]
= vserver services name-service cache unix-user user-by-name delete-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete all the entries for the vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user user-by-name delete-
all` command removes all the user entries that are cached by the user name
for the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the user entries that
are cached by user name need to be deleted.

== Examples

The following example deletes all the cached user entries for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-user user-by-name delete-all -vserver vs0

```
[[IDab2f6dacaf59940d7c6acff3df68be03]]
= vserver services name-service cache unix-user user-by-name delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user user-by-name delete`
command removes the user entries that are cached by the user name. If user
cache propagation is enabled, the corresponding user-by-id cache will also
be removed.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the user entries that
are cached by user name need to be deleted.

``-pw-name <text>`` - pw_name field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the user name for which the cached entries
need to be deleted.

== Examples

The following example deletes all the cached user entries for Vserver vs0
and user name abc:
```

cluster1::> vserver services name-service cache unix-user user-by-name delete -vserver vs0 -pw-name abc

```
[[IDe9e4c787115f5c1f333c513afcef79e7]]
= vserver services name-service cache unix-user user-by-name show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display password struct

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service cache unix-user user-by-name show`
command displays the user information cached by the user name.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
Use this parameter to display detailed information about the user entries
cached by the user name.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the user entries that
are cached by the user name need to be displayed.

``[-pw-name <text>]`` - pw_name field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached user
entries that have the specified user name.

``[-pw-uid <integer>]`` - pw_uid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached user
entries that have the specified user identifier or UID.

``[-pw-gid <integer>]`` - pw_gid field [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to display information only about the cached user
entries that have the specified group identifier or GID.

``[-create-time <MM/DD/YYYY HH:MM:SS>]`` - Create Time [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to display information only about the user entries that
were cached at the specified time.
```

```
``[-source {none|files|dns|nis|ldap|netgrp_byname|dc}]`` - Source of the
Entry [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about the user entries
cached by user name that have the specified look-up source.

== Examples

The following example displays all the users which are cached by user
name:
```

cluster1::> vserver services name-service cache unix-user user-by-name show

```
The following example displays all the users entries cached by user name
for Vserver vs0:
```

cluster1::> vserver services name-service cache unix-user user-by-name show -vserver vs0

```
[[IDf24f3dd22fa484e6f4f599d38d4f1ead]]
= vserver services name-service dns check
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display validation status of a DNS configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service dns check` command to check the
status of configured DNS servers.

== 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.

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver whose DNS mapping needs to be
validated.

``[-name-server <IP Address>]`` - Name Server::
Use this parameter to display information only about name-servers that
match the value you specify.

``[-status {up|down}]`` - Name Server Status::
Use this parameter to display information only about name-servers with a
status that matches the value you specify.

``[-status-details <text>]`` - Status Details::
Use this parameter to display information only about name-servers with
status details that match the value you specify.

== Examples

The following example checks the DNS server mapping on the Vserver vs0:
```

cluster1::> vserver services name-service dns check -vserver vs0
Vserver Name Server Status Status Details
--- --- --- ---
vs0 10.11.12.13 up Response time (msec): 55
vs0 10.11.12.14 up Response time (msec): 70
vs0 10.11.12.15 down Connection refused.

```
[[IDccb6a1d0d8d56dcecb7298366220a7cb]]
= vserver services name-service dns create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new DNS table entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service dns create` command creates new  DNS
server mappings. DNS servers provide remote connection information, such
```

```
as  IP addresses, based on domain and system names.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver on which to create the new DNS
server mapping.

``-domains <text>,...`` - Domains::
Use this parameter to specify the domains of the Vserver. Separate
multiple domains with commas.

``-name-servers <IP Address>,...`` - Name Servers::
Use this parameter to specify the IP addresses of the DNS servers that
provide name service for the domains in this DNS server mapping. Separate
multiple addresses with commas.

``[-timeout <integer>]`` - Timeout (secs)::
Use this parameter to specify a timeout value (in seconds) for queries to
the name servers. The default value is 2 seconds.

``[-attempts <integer>]`` - Maximum Attempts::
Use this parameter to specify the number of attempts the Vserver should
make when querying the DNS name servers. The default value is 1 attempt.

``[-is-tld-query-enabled {true|false}]`` - Is TLD Query Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to enable or disable top-level domain (TLD) queries. If
the parameter is set to ``_false_`` , the resolver will not attempt to
resolve a name that has no "." characters in it. The default value for
this parameter is ``_true_`` .

``[-require-source-address-match {true|false}]`` - Require Source and
Reply IPs to Match [.advaced-priv]#(privilege: advanced)#::
Use this parameter to allow dns responses sourced from an IP that does not
match where the vserver sent the request. If the parameter is set to
``_false_`` , the resolver will allow response from an IP other than the
one to which the request was sent. The default value for this parameter is
``_true_`` .

``[-require-packet-query-match {true|false}]`` - Require Packet Queries to
Match [.advaced-priv]#(privilege: advanced)#::
Use this parameter to check if the query section of the reply packet is
equal to that of the query packet. If the parameter is set to ``_false_``
, the resolver will not check if the query section of the reply packet is
equal to that of the query packet. The default value for this parameter is
```

```
``_true_`` .

``[-skip-config-validation <true>]`` - Skip Configuration Validation::
Use this parameter to skip the DNS configuration validation.
+
The domain name specified with the `-domains` is validated with the
following rules:
+
  * The name must contain only the following characters: A through Z, a
through z, 0 through 9, ".", "-" or "_".
  * The first character of each label, delimited by ".", must be one of
the following characters: A through Z or a through z or 0 through 9.
  * The last character of each label, delimited by ".", must be one of the
following characters: A through Z, a through z, or 0 through 9.
  * The top level domain must contain only the following characters: A
through Z, a through z.
  * The maximum supported length is 255 characters.
  * The system reserves the following names: "all", "local", and
"localhost".



+
The hosts specified with the `-name-servers` parameter are validated to
verify that each of the name servers is reachable, and is providing DNS
services.
+
The validation fails, if the domain name is invalid, or there is no valid
name server.

== Examples

This example creates a new DNS server mapping for the Vserver vs0 in the
domain example.com, specifying that 192.168.0.16 and 192.168.0.24 are the
name servers for this domain.
```

cluster1::> vserver services name-service dns create -vserver vs0 -domains example.com -name-servers 192.168.0.16,192.168.0.24

```
[[ID507f46c4812ea7a087d87fa3bf934ac]]
= vserver services name-service dns delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a DNS table entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service dns delete` command removes the DNS
server mapping from a Vserver.

Deleting a DNS server mapping removes it permanently. If you delete a DNS
server mapping, commands or jobs that do not use IP addresses do not
succeed.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver whose DNS server mapping is
deleted.

== Examples

This example removes the DNS server mapping from the Vserver node1.
```

cluster1::> vserver services name-service dns delete -vserver vs0

```
[[ID45e5c1982a5c3329225a3abbbd55542c]]
= vserver services name-service dns modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change a DNS table entry

*Availability:* This command is available to _cluster_ and _Vserver_
```

```
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service dns modify` command to modify an
existing DNS server mapping.

To permanently remove a mapping, use the xref:{relative_path}vserver-
services-name-service-dns-delete.html[vserver services name-service dns
delete] command.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver whose DNS mapping is modified.

``[-domains <text>,...]`` - Domains::
Use this parameter to specify a domain for the Vserver.

``[-name-servers <IP Address>,...]`` - Name Servers::
Use this parameter to specify the IP addresses of the DNS name servers for
this Vserver.

``[-timeout <integer>]`` - Timeout (secs)::
Use this parameter to specify a timeout value (in seconds) for queries to
the DNS servers.

``[-attempts <integer>]`` - Maximum Attempts::
Use this parameter to specify the number of times to attempt queries to
the DNS servers.

``[-is-tld-query-enabled {true|false}]`` - Is TLD Query Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to enable or disable top-level domain (TLD) queries. If
the parameter is set to ``_false_`` , the resolver will not attempt to
resolve a name that has no "." characters in it. The default value for
this parameter is ``_true_`` .

``[-require-source-address-match {true|false}]`` - Require Source and
Reply IPs to Match [.advaced-priv]#(privilege: advanced)#::
Use this parameter to allow dns responses sourced from an IP that does not
match where the vserver sent the request. If the parameter is set to
``_false_`` , the resolver will allow response from an IP other than the
one to which the request was sent.

``[-require-packet-query-match {true|false}]`` - Require Packet Queries to
Match [.advaced-priv]#(privilege: advanced)#::
```

```
Use this parameter to check if the query section of the reply packet is
equal to that of the query packet. If the parameter is set to ``_false_``
, the resolver will not check if the query section of the reply packet is
equal to that of the query packet.

``[-skip-config-validation <true>]`` - Skip Configuration Validation::
Use this parameter to skip the DNS configuration validation.
+
The domain name specified with the `-domains` is validated with the
following rules:
+
  * The name must contain only the following characters: A through Z, a
through z, 0 through 9, ".", "-" or "_".
  * The first character of each label, delimited by ".", must be one of
the following characters: A through Z or a through z or 0 through 9.
  * The last character of each label, delimited by ".", must be one of the
following characters: A through Z, a through z, or 0 through 9.
  * The top level domain must contain only the following characters: A
through Z, a through z.
  * The maximum supported length is 255 characters.
  * The system reserves the following names: "all", "local", and
"localhost".



+
The hosts specified with the `-name-servers` parameter are validated to
verify that each of the name servers is reachable, and is providing DNS
services.
+
The validation fails, if the domain name is invalid, or there is no valid
name server.

== Examples

This example modifies the DNS server mapping for the domain example.com on
the Vserver vs0, specifying that 10.0.0.1 and 10.0.0.2 are the name
servers for this domain.
```

cluster1::> vserver services name-service dns modify -vserver vs0 -domains example.com -name-servers 10.0.0.1,10.0.0.2

```
== Related Links

* xref:{relative_path}vserver-services-name-service-dns-
delete.html[vserver services name-service dns delete]
```

```
[[IDab4de3447b867845e7e88b1f95827361]]
= vserver services name-service dns show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display DNS configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service dns show` command displays information
about DNS server mappings. DNS servers provide remote connection
information, such as IP addresses, based on domain and system names.

== 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.

``[-vserver <vserver name>]`` - Vserver::
Use this parameter to display information only about the DNS server
mapping of the Vservers you specify.

``[-domains <text>,...]`` - Domains::
Use this parameter to display information only about the DNS server
mappings for Vservers in the domains you specify.

``[-name-servers <IP Address>,...]`` - Name Servers::
Use this parameter to display information only about DNS server mappings
that use the DNS name servers you specify.

``[-timeout <integer>]`` - Timeout (secs)::
Use this parameter to display information only about DNS server mappings
```

```
that have the timeout value you specify.

``[-attempts <integer>]`` - Maximum Attempts::
Use this parameter to display information only about DNS server mappings
that make the maximum number of attempts you specify.

``[-is-tld-query-enabled {true|false}]`` - Is TLD Query Enabled?
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about DNS server mappings
that have the specified TLD query setting.

``[-require-source-address-match {true|false}]`` - Require Source and
Reply IPs to Match [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about DNS server mappings
that have the specified setting to require the source address of the
response packet to match the address where the vserver sent the request.

``[-require-packet-query-match {true|false}]`` - Require Packet Queries to
Match [.advaced-priv]#(privilege: advanced)#::
Use this parameter to display information only about DNS server mappings
that have the specified setting to require the query section of the reply
packet to match that of the query packet.

== Examples

The following example shows typical output from the command. Note that
cluster1 uses different name servers for example.com.
```

cluster1::> vserver services name-service dns show
Name
Vserver Domains Servers
--- --- ---
vs1 example.com 10.0.0.1,
10.0.0.2
vs2 example.com, example2.com 10.0.0.1,
10.0.0.2
vs3 example.com, example2.com 192.168.0.1,
192.168.0.2

```
[[ID3cfb8339c3daf65e543f9e707df8811b]]
= vserver services name-service dns dynamic-update modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Modify a Dynamic DNS Update Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service dns dynamic-update modify` command
modifies the configuration for dynamic DNS updates for a Data Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver for which you want to modify the
dynamic DNS update configuration.

``[-is-enabled {true|false}]`` - Is Dynamic DNS Update Enabled?::
Use this parameter with value `true` to enable the dynamic DNS update
feature. This field is set to `false` by default.

``[-use-secure {true|false}]`` - Use Secure Dynamic Update?::
Use this parameter with value `true` to enable secure dynamic DNS updates.
This field is set to `false` by default.

``[-vserver-fqdn <text>]`` - Vserver FQDN to Be Used for DNS Updates::
Use this parameter to modify the Vserver FQDN to be used for dynamic DNS
updates.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live for DNS Updates [.advaced-priv]#(privilege: advanced)#::
Use this parameter to modify the Time to Live value for the dynamic DNS
updates. The default value is set to 24 hours. The maximum supported value
for TTL is 720 hours.

``[-skip-fqdn-validation <true>]`` - Skip Vserver FQDN Validation::
If the parameter is specified, the FQDN name validation is skipped.

== Examples

The following example enables the dynamic DNS update feature and modifies
the FQDN to be used for dynamic DNS updates for the Vserver vs1,
specifying vs1.abcd.com as the new FQDN.
```

cluster1::*> vserver services name-service dns dynamic-update modify -vserver vs1 -is-enabled true -vserver -fqdn vs1.abcd.com

```
The following example modifies the dynamic DNS updates configuration to
only send secure updates to the DNS server configured for the Vserver vs1.
```

cluster1::*> vserver services name-service dns dynamic-update modify -vserver vs1 -is-enabled true -use -secure true

```
[[ID1f51897521fdbbe76833e2a48e191c63]]
= vserver services name-service dns dynamic-update show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Dynamic DNS Update Configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service dns dynamic-update show` command shows
the dynamic DNS update configuration related to the DNS server for a
Vserver.

== 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.

``[-vserver <vserver name>]`` - Vserver::
Use this parameter to display dynamic DNS update configuration for the
Vservers you specify.

``[-is-enabled {true|false}]`` - Is Dynamic DNS Update Enabled?::
Use this parameter with value `true` to display information about dynamic
DNS update configurations that are active.

``[-use-secure {true|false}]`` - Use Secure Dynamic Update?::
```

```
Use this parameter with value `true` to display information about dynamic
DNS update configurations that are set to send secure dynamic updates
only.

``[-vserver-fqdn <text>]`` - Vserver FQDN to Be Used for DNS Updates::
Use this parameter to display information about dynamic DNS update
configurations that are set to send the dynamic updates with the FQDN you
have specified.

``[-ttl <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Time to
Live for DNS Updates [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays dynamic DNS update
configurations having the specified Time to Live value .

== Examples

The following example shows all information about dynamic DNS update
configurations.
```

cluster1::*> vserver services name-service dns dynamic-update show
gupgclust-3::> dns dynamic-update show
Vserver Is-Enabled Use-Secure Vserver FQDN TTL
--- --- --- --- ---
vs1 true false vs1.abcd.com 24h
vs2 false false vs2.abcd.com 24h
2 entries were displayed.

```
[[IDfd050228e971a33dff7e978e72f7b979]]
= vserver services name-service dns dynamic-update record add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Adds a New DNS Resource Record

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service dns dynamic-update record add` command
sends an update to add a new DNS resource record of an existing logical
interface (LIF) of the Vserver to the configured DNS server.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which you want to add a
resource record on the configured DNS server.

``-lif <lif-name>`` - Logical Interface [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Logical Interface(LIF) name for which
you want to add a resource record on the configured DNS server.

== Examples

The following example adds a resource record entry for the Logical
Interface lif1 belonging to the Vserver vs1 to the configured DNS server.
```

cluster1::*> vserver services name-service dns dynamic-update record add -vserver vs1 -lif lif1

```
[[ID8f46e4f49388395137f080fa63f736fb]]
= vserver services name-service dns dynamic-update record delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Deletes a DNS Resource Record

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service dns dynamic-update record delete`
command sends an update to remove an existing DNS resource record of the
Logical Interface (LIF) of the Vserver from the configured DNS server.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver of which you want to delete a
resource record from the configured DNS server.

{ ``-lif <lif-name>`` - Logical Interface [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Logical Interface(LIF) name whose
corresponding resource record you want to remove from the configured DNS
server.

| ``-address <IP Address>`` - IP Address [.advaced-priv]#(privilege:
advanced)# }::
Use this parameter to specify the IP address of the Logical Interface
whose corresponding resource record you want to remove from the configured
DNS server.

== Examples

The following example removes a resource record entry of the Logical
Interface lif1 belonging to the Vserver vs1 from the configured DNS
server.
```

cluster1::*> vserver services name-service dns dynamic-update record delete -vserver vs1 -lif lif1

```
The following example removes a resource record entry of the Logical
Interface whose address is 1.1.1.1 belonging to the Vserver vs1 from the
configured DNS server.
```

cluster1::*> vserver services name-service dns dynamic-update record delete -address 1.1.1.1 -vserver vs1

```
[[ID418601a0229bfd93d9bbf64117dc9561]]
= vserver services name-service dns hosts create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new host table entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service dns hosts create` command to create
new DNS host table entries. These entries map hostnames to IP addresses.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver on which the host table entry
will be created.

``-address <IP Address>`` - IP Address::
Use this parameter to specify the IP address of the new host table entry.

``-hostname <text>`` - Canonical Hostname::
Use this parameter to specify the full hostname for the new host table
entry.

``[-aliases <text>,...]`` - Aliases::
Use this parameter to specify any aliases to include in the new host table
entry. Separate multiple aliases with commas.

== Examples

This example creates a new DNS host table entry for 10.0.0.17 on the
vserver vs1, with the hostname test.example.com and the alias test.
```

cluster1::> vserver services name-service dns hosts create -vserver vs1 -address 10.0.0.17 -hostname test.example.com -aliases test

```
[[IDf55a51192d8722413c2617f170b3a960]]
= vserver services name-service dns hosts delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a host table entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service dns hosts delete` command to delete
DNS host table entries.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver whose host table entry will be
deleted.

``-address <IP Address>`` - IP Address::
Use this parameter to specify the IP address of the host table entry to
delete.

== Examples

This example removes the DNS host table entry of 10.0.0.15 from the host
table of the vserver vs1.
```

cluster1::> vserver services name-service dns hosts delete -vserver vs1 -address 10.0.0.16

1 entry was deleted.

```
[[ID1b09ae1e80054ad0c85240d6ff41134c]]
= vserver services name-service dns hosts modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify hostname or aliases

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service dns hosts modify` command to modify
existing DNS host table entries.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver whose host table will be
modified.

``-address <IP Address>`` - IP Address::
Use this parameter to specify the IP address of the host table entry to
modify.

``[-hostname <text>]`` - Canonical Hostname::
Use this parameter to specify a full hostname for the host table entry.

``[-aliases <text>,...]`` - Aliases::
Use this parameter to specify alternate hostnames for the host table
entry.

== Examples

This example changes the host table of vserver vs1 so that the hostname
stored in the host table entry for 10.0.0.57 is pgh.example.com.
```

cluster1::> vserver services name-service dns hosts modify -vserver -vs1 -address 10.0.0.57 -hostname pgh.example.com
1 entry was modified.

```
This example changes the host table of vserver vs1 to store the name
loghost as an alternate hostname for IP address 10.0.0.5.
```

cluster1::> vserver services name-service dns hosts modify -vserver vs1 -address 10.0.0.5 -aliases loghost
1 entry was modified.

```
[[ID8042ae808f83b7cb55babb755272401e]]
= vserver services name-service dns hosts show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display IP address to hostname mappings

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service dns hosts show` command to display
Domain Name System (DNS) host table entries. These entries map hostnames
to IP addresses. Entries may also include alternate hostnames, known as
aliases. Host table entries enable you to refer to other Internet hosts by
a memorable name instead of by a numeric IP address. This host table is
similar to the /etc/hosts file found on most UNIX style systems.

== 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.

``[-vserver <vserver name>]`` - Vserver::
Use this parameter to display information only about host table entries on
the Vservers you specify.

``[-address <IP Address>]`` - IP Address::
Use this parameter to display information only about host table entries
```

```
that match the IP addresses you specify.

``[-hostname <text>]`` - Canonical Hostname::
Use this parameter to display information only about host table entries
that match the hostnames you specify.

``[-aliases <text>,...]`` - Aliases::
Use this parameter to display information only about host table entries
that include the alternate hostnames you specify.

== Examples

The following example shows a typical host table.
```

cluster1::> vserver services name-service dns hosts show
Vserver Address Hostname Aliases
---------- ---------- ---------- ----------
vs1 10.0.0.10 mail.example.com
mail, mailhost, snmp
vs1 10.0.0.15 ftp.example.com ftp
vs1 10.0.0.16 www.example.com www
vs2 10.0.0.10 mail.example.com
mail, mailhost, snmp
vs2 10.0.0.15 ftp.example.com ftp
vs2 10.0.0.16 www.example.com www
vs2 10.0.0.17 test.example.com
7 entries were displayed.

```
[[IDdf3ab2fb30fe25b976ac64894de49242]]
= vserver services name-service getxxbyyy getaddrinfo
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the IP address information by using the host name.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy getaddrinfo ` gets the IP
address information by using the host name for a given Vserver. The
underlying service for doing the lookup is selected based on the
configured name service switch order.
```

```
== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-hostname <text>`` - Host Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the Host Name for which the IP address
information is needed

``[-address-family {ipv4|ipv6|all}]`` - Return Addresses for Family
[.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the Address Family for which the IP address
information is needed

``[-show-source {true|false}]`` - Show Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

``[-use-cache <true>]`` - Enable/Disable cache [.advaced-priv]#(privilege:
advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
value is ``_false_`` .

== Examples

The following example requests address information for localhost:
```

cluster1::*> vserver services name-service getxxbyyy getaddrinfo -node cluster1-01 -vserver vs1 -hostname localhost -address-family all -show-source true -use-cache false
Source used for Lookup: Files
Host name: localhost
Canonical name: localhost
IPv4 : 127.0.0.1
IPv6 : ::1

```
[[ID2864c7547fc92ecbaabbc7259c2a8dd9]]
= vserver services name-service getxxbyyy getgrbygid
:hardbreaks:
```

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the group members by using the group identifier or GID.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy getgrbygid ` gets the group
members by using the group identifier or GID for a given Vserver. The
underlying service for doing the lookup is selected based on the
configured name service switch order.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-groupID <integer>`` - Group ID [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the GroupID for which the members are
requested

``[-show-source {true|false}]`` - Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

``[-use-cache <true>]`` - Use Locally-Cached Values [.advaced-
priv]#(privilege: advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
value is ``_false_`` .

== Examples

 The following example requests group information for the given groupid
```

cluster1::*> vserver services name-service getxxbyyy getgrbygid -node cluster1-01 -vserver vs1 -groupID 1
name: daemon
gid: 1

```
[[IDbd2b3bd5bf92f90a137bfaaf187457b5]]
= vserver services name-service getxxbyyy getgrbyname
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the group members by using the group name.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy getgrbyname ` gets the group
members by using the group name.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-groupname <text>`` - Group Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Group Name for which the members are
requested

``[-show-source {true|false}]`` - Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

``[-use-cache <true>]`` - Use Locally-Cached Values [.advaced-
priv]#(privilege: advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
```

```
value is ``_false_`` .

== Examples

The following example requests group information for the given group name
```

cluster1::*> vserver services name-service getxxbyyy getgrbyname -node cluster1-01 -vserver vs1 -groupname daemon -show-source false
name: daemon
gid: 1

```
[[IDf25b5603f7d3f07eba01dfbe27af7ee0]]
= vserver services name-service getxxbyyy getgrlist
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the group list by using the user name.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy getgrlist ` gets the list of
groups to which user belongs. This command will go through all the sources
configured for the group database in the name servers ns-switch
configuration.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-username <text>`` - User Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to retrieve the list of groups where the given user is
a member

``[-use-cache <true>]`` - Use Locally-Cached Values [.advaced-
priv]#(privilege: advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
value is ``_false_`` .

== Examples

The following example requests the grouplist for the given username
```

cluster1::*> vserver services name-service getxxbyyy getgrlist -node cluster1-01 -vserver vs1 -username root

pw_name: root
Groups: 5

```
[[IDa9b6d1ad7e45c719a100ac793b59cba7]]
= vserver services name-service getxxbyyy gethostbyaddr
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the host information from the IP address.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy gethostbyaddr ` gets the
host name by using the IP address. The underlying service for doing the
lookup is selected based on the configured name service switch order.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-ipaddress <IP Address>`` - IP Address [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the IPv4/IPv6 address for which the host
information is needed

``[-show-source {true|false}]`` - Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

``[-use-cache <true>]`` - Enable/Disable cache [.advaced-priv]#(privilege:
advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
```

```
value is ``_false_`` .

== Examples

The following example requests host information for the given IP address:
```

cluster1::*> vserver services name-service getxxbyyy gethostbyaddr -node cluster1-01 -vserver vs1 -ipaddress 127.0.0.1 -show-source false -use-cache false
IP address: 127.0.0.1
Host name: localhost

```
[[ID280950a25438140f095dcc1a8e5ca0fb]]
= vserver services name-service getxxbyyy gethostbyname
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the IP address information from host name.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy gethostbyname ` gets the IP
address by using the host name. The underlying service for doing the
lookup is selected based on the configured name service switch order. When
the look up happens from the hosts file, only the first IP address is
returned for a host configured with multiple IP addresses.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
 Node Use this parameter to specify the node where the lookup will be
performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Vserver Name Use this parameter to specify the Vserver where the lookup
will be performed

``-hostname <text>`` - Host Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the Hostname for which the IP address
information is requested

``[-show-source {true|false}]`` - Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

== Examples

The following example requests IP Address information from the given
hostname
```

cluster1::*> vserver services name-service getxxbyyy gethostbyname -node cluster1-01 -vserver vs1 -hostname localhost -show-source false
Host name: localhost
Canonical name: localhost
IPv4: 127.0.0.1

```
[[IDebccab2394aa7597e0fd1117fb5dface]]
= vserver services name-service getxxbyyy getnameinfo
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the name information by using the IP address.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy getnameinfo ` gets the host
and service by using the socket address. The underlying service for doing
the lookup is selected based on the configured name service switch order.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-ipaddress <IP Address>`` - IP Address [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify IPv4/IPv6 address for which the name
information is requested

``[-show-source {true|false}]`` - Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

``[-use-cache <true>]`` - Enable/Disable cache [.advaced-priv]#(privilege:
```

```
advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
value is ``_false_`` .

== Examples

The following example gets the name information for the given IP Address:
```

cluster1::*> vserver services name-service getxxbyyy getnameinfo -node cluster1-01 -vserver vs1 -ipaddress 127.0.0.1 -show-source false -use-cache false
IP address: 127.0.0.1
Host name: localhost

```
[[ID03235ac61d40366e863c1ef91b2cd634]]
= vserver services name-service getxxbyyy getpwbyname
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the password entry by using the user name.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy getpwbyname ` gets the
password entry by using the user name. The underlying service for doing
the lookup is selected based on the configured name service switch order.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-username <text>`` - User Name [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the Username for which the password entry is
requested
```

```
``[-show-source {true|false}]`` - Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

``[-use-cache <true>]`` - Enable/Disable cache [.advaced-priv]#(privilege:
advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
value is ``_false_`` .

== Examples

The following example requests password entry from the given username:
```

cluster1::*> vserver services name-service getxxbyyy getpwbyname -node cluster1-01 -vserver vs1 -username vsadmin -show-source true -use-rbac false -use-cache false
Source used for lookup: Files
pw_name: daemon
pw_passwd: *
pw_uid: 1, pw_gid: 1
pw_gecos: Owner of many system processes
pw_dir: /root
pw_shell: /usr/sbin/nologin

```
[[IDf2b9ea0baaf3009c9820330786a22206]]
= vserver services name-service getxxbyyy getpwbyuid
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Gets the password entry by using the user identifier or UID.

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy getpwbyuid ` gets the
password entry by using the user identifier or UID. The underlying service
for doing the lookup is selected based on the configured name service
switch order.

== Parameters
```

```
``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-userID <integer>`` - User ID [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the UserID for whom the password entry is
requested

``[-show-source {true|false}]`` - Source used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

``[-use-cache <true>]`` - Enable/Disable cache [.advaced-priv]#(privilege:
advanced)#::
If set to ``_true_`` , locally-cached values will be used. The default
value is ``_false_`` .

== Examples

The following example requests password entry by using the user ID:
```

cluster1::*> vserver services name-service getxxbyyy getpwbyuid -node cluster1-01 -vserver vs1 -userID 1001 -show-source true -use-rbac true -use-cache false
Source used for Lookup: Files
pw_name: vsadmin
pw_passwd: 1f7b22f68$KihT1ptYqpEjcM4jfE60f0
pw_uid: 1001
pw_gid: 65533
pw_gecos: User
pw_dir: /var/home/vsadmin
pw_shell: /sbin/ngsh

```
[[ID4bcea87a882c95afa91e04c409ec1a8f]]
= vserver services name-service getxxbyyy netgrpcheck
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Check if a client is part of a netgroup using combined API
```

```
*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The ` vserver services name-service getxxbyyy netgrpcheck ` checks if a
client is part of a netgroup. The underlying service for doing the lookup
is selected based on the configured name service switch order.

== Parameters

``-node {<nodename>|local}`` - Node Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the node where the lookup will be performed

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver where the lookup will be
performed

``-netgroup <text>`` - Netgroup Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Netgroup name

``-clientIP <IP Address>`` - Client IP Address [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Client IP for which the membership in a
given netgroup needs to be checked

``[-enable-domain-search-flag {true|false}]`` - Use DNS domain [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to use DNS domain. Default value for this field is true

``[-trust-any-source {true|false}]`` - Trust any source [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to set trust any source parameter. Default value for
this field is false

``[-show-source {true|false}]`` - Source Used for Lookup [.advaced-
priv]#(privilege: advanced)#::
Use this parameter to specify if source used for lookup needs to be
displayed

== Examples

The following example checks if the given client is part of the given
```

```
netgroup:
```

cluster1::*> vserver services name-service getxxbyyy netgrpcheck -node cluster1-01 -vserver vs1 -netgroup net1 -clientIP 10.232.98.198 -show-source false
10.232.98.198 is a member of net1

```
[[ID1c4d2dd97566883b41e148221a78df99]]
= vserver services name-service ldap check-ipv6
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display validation status of an IPv6 LDAP configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service ldap check-ipv6` command to check
the status of an IPv6 LDAP configuration.

== 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.

``-vserver <Vserver Name>`` - Vserver::
Use this parameter to specify the Vserver whose LDAP configuration needs
to be validated.

``[-client-config <text>]`` - Client Configuration Name::
Use this parameter to display information only about LDAP client
configuration which is assigned to LDAP configuration for the specified
Vserver.

``[-ldap-status {up|down}]`` - LDAP Status::
Use this parameter to display information only about LDAP configurations
```

```
with a status that matches the value you specify.

``[-ldap-status-details <text>]`` - LDAP Status Details::
 Use this parameter to display information only about LDAP configurations
with a status detail that matches the value you specify.

``[-ldap-dn-status-details <text>,...]`` - LDAP DN Status Details::
 Use this parameter to display information only about LDAP DN
configurations with a status detail that matches the value you specify.

== Examples

The following examples check the LDAP configuration on the SVM vs1:
```

cluster1::> vserver services name-service ldap check-ipv6 -vserver vs1
Vserver: vs1
Client Configuration Name: clientcfg1
LDAP Status: up
LDAP Status Details: Successfully connected to LDAP server "2001:db8:3333:4444:5555:6666:7777:8888".
LDAP DN Status Details: All the configured DNs are available.

cluster1::> vserver services name-service ldap check-ipv6 -vserver vs1
Vserver: vs1
Client Configuration Name: clientcfg1
LDAP Status: up
LDAP Status Details: Successfully connected to LDAP server "2001:db8:3333:4444:5555:6666:7777:8888".
LDAP DN Status Details: Validation of Domains specified in the LDAP client configuration failed. Reason: bind-dn is invalid or bind credentials are invalid. Correct the configuration and try again.

```
In the above example, you can correct the LDAP configuration by performing
either of the following procedures: -- If the bind-dn is invalid, use the
xref:{relative_path}vserver-services-name-service-ldap-client-
modify.html[vserver services name-service ldap client modify] command to
correct it. -- If the bind credentials are invalid, use the
xref:{relative_path}vserver-services-name-service-ldap-client-modify-bind-
password.html[vserver services name-service ldap client modify-bind-
password] command to correct them.

== Related Links

* xref:{relative_path}vserver-services-name-service-ldap-client-
modify.html[vserver services name-service ldap client modify]
* xref:{relative_path}vserver-services-name-service-ldap-client-modify-
bind-password.html[vserver services name-service ldap client modify-bind-
password]
```

```
[[ID6aebf68e09fe16155ed4cac4c8a1edd4]]
= vserver services name-service ldap check
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display validation status of a LDAP configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service ldap check` command to check the
status of the LDAP configuration.

== 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.

``-vserver <Vserver Name>`` - Vserver::
Use this parameter to specify the Vserver whose LDAP configuration needs
to be validated.

``[-client-config <text>]`` - Client Configuration Name::
Use this parameter to specify the LDAP client configuration which is
assigned to LDAP configuration for the specified Vserver.

``[-ldap-status {up|down}]`` - LDAP Status::
Use this parameter to display information only about LDAP configurations
with a status that matches the value you specify.

``[-ldap-status-details <text>]`` - LDAP Status Details::
 Use this parameter to display information only about LDAP configurations
```

```
with a status detail that matches the value you specify.

``[-ldap-dn-status-details <text>,...]`` - LDAP DN Status Details::
 Use this parameter to display information only about LDAP DN
configurations with a status detail that matches the value you specify.

== Examples

The following examples check the LDAP configuration on the SVM vs0:
```

cluster1::> vserver services name-service ldap check -vserver vs0
Vserver: vs0
Client Configuration Name: c1
LDAP Status: up
LDAP Status Details: Successfully connected to LDAP server "10.11.12.13".
LDAP DN Status Details: All the configured DNs are available.

```

```

cluster1::> vserver services name-service ldap check -vserver vs0
Vserver: vs0
Client Configuration Name: c1
LDAP Status: up
LDAP Status Details: Successfully connected to LDAP server "10.11.12.13".
LDAP DN Status Details: Validation of Domains specified in the LDAP client configuration failed. Reason: bind-dn is invalid or bind credentials are invalid. Correct the configuration and try again.
In the above example, you can correct the LDAP configuration by performing either of the following procedures:
— If the bind-dn is invalid, use the "ldap client modify" command to correct it.
— If the bind credentials are invalid, use the "ldap client modify-bind-password" command to correct them.

```
[[ID6d540389ce7f2985dfbb0593e1149f1b]]
= vserver services name-service ldap create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an LDAP configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap create` command associates an LDAP
client configuration with a Vserver.
```

```
== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver with which you want to associate the
LDAP client configuration. A data Vserver or admin Vserver can be
specified.

``-client-config <text>`` - LDAP Client Configuration::
This parameter specifies the name of the LDAP client configuration,
defined under the `vserver services name-service ldap client` command,
that you want to associate with the Vserver. The value of the bind-as-
cifs-server parameter on this LDAP client should be false, if the CIFS
server of the associated data Vserver does not exist or exists in
workgroup mode.

``[-skip-config-validation <true>]`` - Skip Configuration Validation::
Use this parameter to skip the LDAP configuration validation.
+
The LDAP client configuration, specified by the `-client-config`
parameter, that you want to associate with the Vserver is validated to
verify that at least one of the LDAP servers is reachable, and is
providing LDAP services.
+
.
+
The validation fails if ONTAP was unable to connect to any LDAP server
with the specified `-client-config` .

== Examples

The following example associates the LDAP client configuration "corp" with
the Vserver "vs1":
```

cluster1::> vserver services name-service ldap create -vserver vs1 -client-config corp

```
[[ID6bb393329eb01645037336028862dd8b]]
= vserver services name-service ldap delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an LDAP configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap delete` command removes the LDAP
configuration, which is an LDAP client configuration's association with a
Vserver.

NOTE: Make sure that you remove 'ldap' from the Vserver's `-ns-switch` and
`-nm-switch` parameters and test connectivity before deleting a working
LDAP configuration.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver from which you want to disassociate
the LDAP client configuration. A data Vserver or admin Vserver can be
specified.

== Examples

The following example disassociates the current LDAP client configuration
from Vserver "vs1".
```

cluster1::> vserver services name-service ldap delete -vserver vs1

```
[[ID5b7ac353ac03a7fef9dac9c7b651a76e]]
= vserver services name-service ldap modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Modify an LDAP configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap modify` command modifies an LDAP
client configuration's association with a Vserver.

NOTE: Make sure that you remove 'ldap' from the Vserver's `-ns-switch` and
`-nm-switch` configurations and test connectivity before disabling a
working LDAP configuration.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver with which you want to associate the
LDAP client configuration. A data Vserver or admin Vserver can be
specified.

``[-client-config <text>]`` - LDAP Client Configuration::
This parameter specifies the name of the LDAP client configuration,
defined under `vserver services name-service ldap client` command, that
you want to associate with the Vserver. The value of the bind-as-cifs-
server parameter on this LDAP client should be false if the CIFS server of
the associated data Vserver does not exist or exists in workgroup mode.

``[-skip-config-validation <true>]`` - Skip Configuration Validation::
Use this parameter to skip the LDAP configuration validation.
+
The LDAP client configuration, specified by the `-client-config`
parameter, that you want to associate with the Vserver is validated to
verify that at least one of the LDAP servers is reachable, and is
providing LDAP services.
+
.
+
The validation fails if ONTAP was unable to connect to any LDAP server
with the specified `-client-config` .

== Examples
```

```
The following example modifies the LDAP client configuration used by
Vserver "vs1" to "corpnew":
```

cluster1::> vserver services name-service ldap modify -vserver vs1 -client-config corpnew

```
[[IDb2d207c79e8921397c3288e8b433beaf]]
= vserver services name-service ldap show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display LDAP configurations

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap show` command displays information
about LDAP configurations.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays information about the
LDAP configuration on the specified Vserver. A data Vserver or admin
Vserver can be specified.

``[-client-config <text>]`` - LDAP Client Configuration::
If you specify this parameter, the command displays information about LDAP
configurations using the specified client.

== Examples

The following example shows the LDAP configuration for Vserver "vs1":
```

cluster1::> vserver services name-service ldap show -vserver vs1
Client
Vserver Configuration
--- ---

```
[[ID0d08ca41545bb09ab3ce077711042730]]
= vserver services name-service ldap client create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an LDAP client configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap client create` command creates an
LDAP client configuration. A client configuration is associated with a
Vserver using the `vserver services name-service ldap` commands.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the Vserver for which configuration is created. A
data Vserver or admin Vserver can be specified.

``-client-config <text>`` - Client Configuration Name::
This parameter specifies the name that you would like to use to refer to
the new LDAP client configuration.

{ ``-ldap-servers <text>,...`` - LDAP Server List::
This parameter specifies the list of LDAP servers used when making LDAP
connections using this client configuration. If you specify this
parameter, you cannot specify the `-servers` , `-ad-domain` or `-
preferred-ad-servers` parameters. This parameter takes both FQDNs and IP
addresses.

| ``-servers <IP Address>,...`` - (DEPRECATED)-LDAP Server List::
(DEPRECATED)This parameter specifies the list of LDAP servers used when
making LDAP connections using this client configuration. If you specify
this parameter, you cannot specify the `-ldap-servers` , `-ad-domain` , `-
preferred-ad-servers` or `-bind-as-cifs-server` parameters. This parameter
is deprecated 9.1.0 and later releases. Use `-ldap-servers` instead.

| ``-ad-domain <TextNoCase>`` - Active Directory Domain::
```

```
This parameter specifies the name of the Active Directory domain used to
discover LDAP servers for use by this client. This assumes that the Active
Directory schema has been extended to act as a NIS replacement. If you use
this parameter, you cannot specify the `-ldap-servers` and `-servers`
parameter. However, you can specify a list of preferred servers using the
`-preferred-ad-servers` parameter.

``[-preferred-ad-servers <IP Address>,...]`` - Preferred Active Directory
Servers::
This parameter specifies a list of LDAP servers that are preferred over
those that are discovered in the domain specified in the `-ad-domain`
parameter.

``[-restrict-discovery-to-site {true|false}]`` - Restrict discovery to
site scope }::
This parameter specifies whether to restrict LDAP server discovery to
site-scope only. The default value is ``_false_`` . The restriction only
applies when the `-ad-domain` parameter is specified as part of the create
command. This can be enabled only if `-default-site` parameter is
specified in the CIFS server configuration.

``[-bind-as-cifs-server {true|false}]`` - Bind Using the Vserver's CIFS
Credentials::
This parameter specifies whether LDAP binds made using this client
configuration use the Vserver's CIFS server credentials. If you do not
specify this parameter, and `-ad-domain` is configured, the default is
`true` , otherwise the default is `false` . Note that the LDAP client
always uses only sasl bind, if `-bind-as-cifs-server` is set to `true` .
The `-min-bind-level` parameter is ignored in this case.

``-schema <text>`` - Schema Template::
This parameter specifies the name of the schema template the Vserver uses
when making LDAP queries. You can view and modify the templates using the
`vserver services name-service ldap client schema` commands.

``[-port <integer>]`` - LDAP Server Port::
This parameter specifies the port the LDAP client uses to connect to LDAP
servers. Default value for port is `636` , if `-ldaps-enabled` parameter
is specified as `true` . Otherwise, default value for port is `389` .

``[-query-timeout <integer>]`` - Query Timeout (sec)::
This parameter specifies the amount of time (in seconds) that the LDAP
client waits for a query to complete. If you do not specify this
parameter, the default is `3` seconds.

``[-min-bind-level {anonymous|simple|sasl}]`` - Minimum Bind
```

```
Authentication Level::
This parameter specifies the lowest acceptable level of security the LDAP
client uses to bind to an LDAP server. If you do not specify this
parameter, the default is `sasl` bind in case `-ad-domain` is configured,
`simple` bind in case `-bind-dn` is configured, otherwise `anonymous`
bind. Note that regardless of the `-min-bind-level` configured, LDAP
client would always start bind mechanism in the order of `sasl` , then
`simple` and lastly `anonymous` . Also, if `-bind-as-cifs-server` is set,
then `-min-bind-level` is ignored, and only `sasl` will be used.

``[-bind-dn <ldap_dn>]`` - Bind DN (User)::
This parameter specifies the user that binds to the LDAP servers. For
Active Directory servers, specify the user in the account (DOMAIN\user) or
principal (user@domain.com) form. Otherwise, specify the user in
distinguished name form, like "CN=user,DC=domain,DC=com" or
"CN=administrator,CN=users,DC=domain,DC=com". This parameter is ignored if
`-bind-as-cifs-server` is set.

``[-base-dn <ldap_dn>]`` - Base DN::
This parameter specifies the default base DN for all searches, including
user, group, and netgroup searches. For example, "DC=example,DC=com". If
you do not specify this parameter, the default is the root, specified by
an empty (`""` ) set.

``[-base-scope {base|onelevel|subtree}]`` - Base Search Scope::
This parameter specifies the default search scope for LDAP queries.
Specify `base` to search just the named entry, `onelevel` to search
entries immediately below the DN, or `subtree` to search the named DN
entry and the entire subtree below the DN. If you do not specify this
parameter, the scope is set to `subtree` by default.

``[-user-dn <ldap_dn>]`` - User DN [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the user DN, which overrides the base DN for user
lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple user or group DNs and a DN
contains a semicolon, add an escape character (\) immediately before the
semicolon or enclose the entire DN with quotation marks (").
+


``[-user-scope {base|onelevel|subtree}]`` - User Search Scope [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the user search scope. If you do not specify this
```

```
parameter, the scope is set to `subtree` by default.

``[-group-dn <ldap_dn>]`` - Group DN [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the group DN, which overrides the base DN for
group lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple user or group DNs and a DN
contains a semicolon, add an escape character (\) immediately before the
semicolon or enclose the entire DN with quotation marks (").
+


``[-group-scope {base|onelevel|subtree}]`` - Group Search Scope [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the group search scope. If you do not specify
this parameter, the scope is set to `subtree` by default.

``[-netgroup-dn <ldap_dn>]`` - Netgroup DN [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the netgroup DN, which overrides the base DN
netgroup lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple netgroup DNs and a DN contains a
semicolon, add an escape character (\) immediately before the semicolon or
enclose the entire DN with quotation marks (").
+


``[-netgroup-scope {base|onelevel|subtree}]`` - Netgroup Search Scope
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the netgroup search scope. If you do not specify
this parameter, the scope is set to `subtree` by default.

``[-use-start-tls {true|false}]`` - Use start-tls Over LDAP Connections::
This parameter specifies whether or not to use Start TLS over LDAP
connections. When enabled, the communication between the Data ONTAP LDAP
Client and the LDAP Server will be encrypted using Start TLS. Start TLS is
a mechanism to provide secure communication by using the TLS/SSL
protocols. If you do not specify this parameter, the default is `false` .

``[-is-netgroup-byhost-enabled {true|false}]`` - Enable Netgroup-By-Host
Lookup [.advaced-priv]#(privilege: advanced)#::
Use this parameter to enable or disable netgroup-by-host lookup. If your
```

```
LDAP directory contains map structures equivalent to the netgroup.byhost
map in NIS, enabling this feature greatly speeds up netgroup resolution
queries over LDAP. By default this parameter is set to false.

``[-netgroup-byhost-dn <ldap_dn>]`` - Netgroup-By-Host DN [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the netgroup-by-host DN, which overrides the base
DN for netgroup-by-host lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple netgroup DNs and a DN contains a
semicolon, add an escape character (\) immediately before the semicolon or
enclose the entire DN with quotation marks (").
+


``[-netgroup-byhost-scope {base|onelevel|subtree}]`` - Netgroup-By-Host
Scope [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the netgroup-by-host search scope for LDAP
queries. If you do not specify this parameter, the scope is set to
`subtree` by default.

``[-session-security {none|sign|seal}]`` - Client Session Security::
This parameter specifies the level of security to be used for LDAP
communications. If you do not specify this parameter, the default is
`none` .
+
LDAP Client Session Security can be one of the following:
+

+
  *  none - No Signing or Sealing.
  *  sign - Sign LDAP traffic.
  *  seal - Seal and Sign LDAP traffic.



``[-referral-enabled {true|false}]`` - LDAP Referral Chasing::
This parameter specifies whether LDAP referral is enabled or not.

``[-group-membership-filter <text>]`` - Group Membership Filter [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the custom LDAP search filter to be used when
looking up group membership from an LDAP server. Examples of valid filters
are "(cn=*99)", "(cn=1*)", "(|(cn=*22)(cn=*33))".

``[-ldaps-enabled {true|false}]`` - Is LDAPS Enabled::
```

```
This parameter specifies whether or not to use LDAPS over LDAP
connections. If you do not specify this parameter, the value will be based
on `port` . If `port` is mentioned as `636` , then the value will be
`true` , otherwise the value will be `false` .

``[-try-channel-binding {true|false}]`` - Try Channel Binding::
This parameter specifies whether channel binding will be tried in case of
LDAP connections to the LDAP server. If you do not specify this parameter,
the default is ``_true_`` . Channel binding will be tried only if `-use-
start-tls` or `-ldaps-enabled` is enabled along with `-session-security`
set to either ``_sign_`` or ``_seal_`` .

== Examples

The following example creates an LDAP client configuration named `corp`
that makes anonymous binds to `ldapserver.example.com` for Vserver `vs1` :
```

cluster1::> vserver services name-service ldap client create -vserver vs1 -client-config corp -ldap-servers ldapserver.example.com

```
The following example creates an LDAP client configuration named `corp`
that makes binds to `ldapserver.example.com` for Vserver `vs1` for bind-dn
`diag` :
```

cluster1::> vserver services name-service ldap client create -vserver vs1 -client-config corp -ldap-servers ldapserver.example.com -bind-dn diag
Please enter password:
Confirm password:

```
The following example creates an LDAP client configuration with multiple
user DNs.

NOTE: The following commands are only available in advanced mode.
```

cluster1::*> vserver services ldap client create -vserver vs1 -client-config corp -ldap-servers ldapserver.example.com
-user-dn "ou=People,dc=mypc,dc=example,dc=com; ou=People1,dc=mypc1,dc=example2,dc=com"

```
The following example creates an LDAP client configuration with multiple
user DNs, one of them containing a semicolon
```

cluster1::*> vserver services ldap client create -vserver vs1 -client-config corp -ldap-servers ldapserver.example.com
-user-dn "ou=People,dc=mypc,dc=example,dc=com; ou=People1,dc=mypc1,dc=example2,dc=com"

```
The following example creates an LDAP client configuration with multiple
user DNs, one of them containing a semicolon and a backslash.
```

cluster1::*> vserver services ldap client create -vserver vs1 -client-config corp -ldap-servers ldapserver.example.com
-user-dn "ou=People\;,dc=mypc,dc=example,dc=com\\; ou=People1,dc=mypc1,dc=example2,dc=com"

```
The following example creates an LDAP client configuration with netgroup
by host DN.
```

cluster1::*>vserver services ldap client create -vserver vs1 -client-config corp -ldap-servers ldapserver.example.com
-netgroup-byhost-dn nisMapName="netgroup.byhost",dc=rfcbis,dc=com

```
The following example creates an LDAP client configuration with ldap-
servers as list of ip addresses.
```

cluster1::*>vserver services ldap client create -vserver vs1 -client-config corp -ldap-servers 172.16.0.100,172.16.0.101
-netgroup-byhost-dn nisMapName="netgroup.byhost",dc=rfcbis,dc=com

```
The following example creates an LDAP client configuration with ldap-
servers as list of ip addresses and hostnames.
```

cluster1::*>vserver services ldap client create -vserver vs1 -client-config corp -ldap-servers ldapserver.example.com,172.16.0.100,172.16.0.101 -netgroup-byhost-dn nisMapName="netgroup.byhost",dc=rfcbis,dc=com

```
[[IDa27ebb6170e6c0703c104c36dc51c55e]]
= vserver services name-service ldap client delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an LDAP client configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap client delete` command deletes an
LDAP client configuration. A Vserver administrator can only delete
configurations owned by the Vserver.

== Parameters

``[-vserver <Vserver Name>]`` - Vserver::
This parameter specifies the name of the Vserver which owns the LDAP
client you want to delete. A data Vserver or admin Vserver can be
specified.

``-client-config <text>`` - Client Configuration Name::
This parameter specifies the name of the LDAP client configuration you
want to delete.

== Examples

The following example deletes an LDAP client configuration named `corp`
owned by Vserver `vs1` :
```

cluster1::> vserver services name-service ldap client delete -vserver vs1 -client-config corp

```
[[ID8642d251c658682cacf3b36464cf9dec]]
= vserver services name-service ldap client modify-bind-password
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify Bind Password of an LDAP client configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The ` vserver services name-service ldap client modify-bind-password `
command modifies bind-password of a given LDAP client configuration.

== Parameters

``[-vserver <Vserver Name>]`` - Vserver::
This parameter specifies the name of the Vserver which owns the LDAP
client you want to modify. A data Vserver or admin Vserver can be
specified.

``-client-config <text>`` - Client Configuration Name::
This parameter specifies the name of the LDAP client configuration.

== Examples

The following example modifies the password for a given LDAP client
configuration
```

cluster1::> vserver services name-service ldap client modify-bind-password -client-config corp
Please enter password:
Confirm password:

```
[[ID45622bb50b1f6edee042660c53021f2a]]
= vserver services name-service ldap client modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Modify an LDAP client configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap client modify` command modifies an
LDAP client configuration. A Vserver administrator can modify only
configurations owned by the Vserver. Use the xref:{relative_path}vserver-
services-name-service-ldap-client-modify-bind-password.html[vserver
services name-service ldap client modify-bind-password] command to modify
the bind password.

== Parameters

``[-vserver <Vserver Name>]`` - Vserver::
This parameter specifies the name of the Vserver which owns the LDAP
client you want to modify. A data Vserver or admin Vserver can be
specified.

``-client-config <text>`` - Client Configuration Name::
This parameter specifies the name of the LDAP client configuration.

{ ``[-ldap-servers <text>,...]`` - LDAP Server List::
This parameter specifies the list of LDAP servers used when making LDAP
connections using this client configuration. If you specify this
parameter, you cannot specify the `-servers` , `-ad-domain` or `-
preferred-ad-servers` parameters.

| ``[-servers <IP Address>,...]`` - (DEPRECATED)-LDAP Server List::
(DEPRECATED)This parameter specifies the list of LDAP servers used when
making LDAP connections using this client configuration. If you specify
this parameter, you cannot specify the `-ldap-servers` , `-ad-domain` , `-
preferred-ad-servers` or `-bind-as-cifs-server` parameters. This parameter
is deprecated 9.1.0 and later releases. Use `-ldap-servers` instead.

| ``[-ad-domain <TextNoCase>]`` - Active Directory Domain::
This parameter specifies the name of the Active Directory domain used to
discover LDAP servers for use by this client. This assumes that the Active
Directory schema has been extended to act as a NIS replacement. If you use
this parameter, you cannot specify the `-servers` , `-ldap-servers`
parameter. However, you can specify a list of preferred servers using the
`-preferred-ad-servers` parameter.

``[-preferred-ad-servers <IP Address>,...]`` - Preferred Active Directory
Servers::
```

```
This parameter specifies a list of LDAP servers that are preferred over
those that are discovered in the domain specified in the `-ad-domain`
parameter.

``[-restrict-discovery-to-site {true|false}]`` - Restrict discovery to
site scope }::
This parameter specifies whether to restrict server discovery to site-
scope only. The default value is ``_false_`` . The restriction only
applies when an `-ad-domain` is configured. This can be enabled only if `-
default-site` parameter is specified in the CIFS server configuration.

``[-bind-as-cifs-server {true|false}]`` - Bind Using the Vserver's CIFS
Credentials::
This parameter specifies whether or not LDAP binds made using this client
configuration use the Vserver's CIFS server credentials. Note that the
LDAP client always uses only sasl bind, if `-bind-as-cifs-server` is set
to `true` . The `-min-bind-level` parameter is ignored in this case.

``[-schema <text>]`` - Schema Template::
This parameter specifies the name of the schema template the Vserver uses
when making LDAP queries. You can view and modify the templates using the
`vserver services name-service ldap client schema` commands.

``[-port <integer>]`` - LDAP Server Port::
This parameter specifies the port the LDAP client uses to connect to LDAP
servers. Default value for port is `636` , if `-ldaps-enabled` parameter
is specified as `true` . Otherwise, default value for port is `389` .

``[-query-timeout <integer>]`` - Query Timeout (sec)::
This parameter specifies the amount of time (in seconds) that the LDAP
client waits for a query to complete. If you do not specify this
parameter, the default is `3` seconds.

``[-min-bind-level {anonymous|simple|sasl}]`` - Minimum Bind
Authentication Level::
This parameter specifies the lowest acceptable level of security the LDAP
client uses to bind to an LDAP server. Note that regardless of the `-min-
bind-level` configured, LDAP client would always start bind mechanism in
the order of `sasl` , then `simple` and lastly `anonymous` . Also, if `-
bind-as-cifs-server` is set, then `-min-bind-level` is ignored, and only
`sasl` will be used.

``[-bind-dn <ldap_dn>]`` - Bind DN (User)::
This parameter specifies the user that binds to the LDAP servers. For
Active Directory servers, specify the user in the account (DOMAIN\user) or
principal (user@domain.com) form. Otherwise, specify the user in
```

```
distinguished name form, like "CN=user,DC=domain,DC=com" or
"CN=administrator,CN=users,DC=domain,DC=com". This parameter is ignored if
`-bind-as-cifs-server` is set.

``[-base-dn <ldap_dn>]`` - Base DN::
This parameter specifies the default base DN for all searches, including
user, group, and netgroup searches. For example, "DC=example,DC=com". If
you do not specify this parameter, the default is the root, specified by
an empty (`""` ) set.

``[-base-scope {base|onelevel|subtree}]`` - Base Search Scope::
This parameter specifies the default search scope for LDAP queries.
Specify `base` to search just the named entry, `onelevel` to search
entries immediately below the DN, or `subtree` to search the named DN
entry and the entire subtree below the DN. If you do not specify this
parameter, the scope is set to `subtree` by default.

``[-user-dn <ldap_dn>]`` - User DN [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the user DN, which overrides the base DN for user
lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple user or group DNs and a DN
contains a semicolon, add an escape character (\) immediately before the
semicolon or enclose the entire DN with quotation marks (").
+


``[-user-scope {base|onelevel|subtree}]`` - User Search Scope [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the user search scope. If you do not specify this
parameter, the scope is set to `subtree` by default.

``[-group-dn <ldap_dn>]`` - Group DN [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the group DN, which overrides the base DN for
group lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple user or group DNs and a DN
contains a semicolon, add an escape character (\) immediately before the
semicolon or enclose the entire DN with quotation marks (").
+
```

```
``[-group-scope {base|onelevel|subtree}]`` - Group Search Scope [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the group search scope. If you do not specify
this parameter, the scope is set to `subtree` by default.

``[-netgroup-dn <ldap_dn>]`` - Netgroup DN [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the netgroup DN, which overrides the base DN
netgroup lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple netgroup DNs and a DN contains a
semicolon, add an escape character (\) immediately before the semicolon or
enclose the entire DN with quotation marks (").
+



``[-netgroup-scope {base|onelevel|subtree}]`` - Netgroup Search Scope
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the netgroup search scope. If you do not specify
this parameter, the scope is set to `subtree` by default.

``[-use-start-tls {true|false}]`` - Use start-tls Over LDAP Connections::
This parameter specifies whether or not to use Start TLS over LDAP
connections. When enabled, the communication between the Data ONTAP LDAP
Client and the LDAP Server will be encrypted using Start TLS. Start TLS is
a mechanism to provide secure communication by using the TLS/SSL
protocols. If you do not specify this parameter, the default is `false` .

``[-is-netgroup-byhost-enabled {true|false}]`` - Enable Netgroup-By-Host
Lookup [.advaced-priv]#(privilege: advanced)#::
Use this parameter to enable or disable netgroup-by-host lookup. If your
LDAP directory contains map structures equivalent to the netgroup.byhost
map in NIS, enabling this feature greatly speeds up netgroup resolution
over LDAP. By default this parameter is set to false.

``[-netgroup-byhost-dn <ldap_dn>]`` - Netgroup-By-Host DN [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the netgroup-by-host DN, which overrides the base
DN for netgroup-by-host lookups.
+
NOTE: To specify multiple DNs, separate multiple DN entries with
semicolons (;). If you configure multiple netgroup DNs and a DN contains a
semicolon, add an escape character (\) immediately before the semicolon or
enclose the entire DN with quotation marks (").
+
```

```
``[-netgroup-byhost-scope {base|onelevel|subtree}]`` - Netgroup-By-Host
Scope [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the netgroup-by-host search scope for LDAP
queries. If you do not specify this parameter, the scope is set to
`subtree` by default.

``[-session-security {none|sign|seal}]`` - Client Session Security::
This parameter specifies the level of security to be used for LDAP
communications. If you do not specify this parameter, the default is
`none` .
+
LDAP Client Session Security can be one of the following:
+

+
  *  none - No Signing or Sealing.
  *  sign - Sign LDAP traffic.
  *  seal - Seal and Sign LDAP traffic.



``[-skip-config-validation <true>]`` - Skip Configuration Validation::
Use this parameter to skip the LDAP client configuration validation.
+
The LDAP client configuration specified with the `-client-config`
parameter is validated to verify that all the Vservers associated with
this LDAP client configuration has at least one of the LDAP servers
reachable, and is providing LDAP services.
+
The validation fails if ONTAP was unable to connect to any LDAP server
with the specified `-client-config` .

``[-referral-enabled {true|false}]`` - LDAP Referral Chasing::
This parameter specifies whether LDAP referral is enabled or not.

``[-group-membership-filter <text>]`` - Group Membership Filter [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the custom LDAP search filter to be used when
looking up group membership from an LDAP server. Examples of valid filters
are "(cn=*99)", "(cn=1*)", "(|(cn=*22)(cn=*33))".

``[-ldaps-enabled {true|false}]`` - Is LDAPS Enabled::
This parameter specifies whether or not to use LDAPS over LDAP
connections. If you do not specify this parameter, the value will be based
on `port` . If `port` is mentioned as `636` , then the value will be
```

```
`true` , otherwise the value will be `false` .

``[-try-channel-binding {true|false}]`` - Try Channel Binding::
This parameter specifies whether to use channel binding for LDAP
connections to the LDAP server. If you do not specify this parameter, the
default is ``_true_`` . Channel binding will be tried only if `-use-start-
tls` or `-ldaps-enabled` is enabled along with `-session-security` set to
either ``_sign_`` or ``_seal_`` .

== Examples

The following example modifies an existing LDAP client configuration named
`corp` owned by Vserver `vs1` to require simple binds using the
administrator@example.com account:
```

cluster1::> vserver services name-service ldap client modify -client-config corp -vserver vs1 -bind-dn administrator@example.com -min-bind-level simple

```
The following example modifies the user DN of an existing LDAP client
configuration to contain multiple DNs separated by a semicolon.
```

cluster1::> vserver services ldap client modify -client-config corp -vserver vs1 -bind-dn administrator@example.com
-user-dn "ou=People,dc=mypc,dc=example,dc=in; ou=People1,dc=mypc,dc=example2,dc=com" -min-bind -level simple

```
The following example demonstrates how you can use a semicolon as a valid
character in a DN instead of a separator.
```

cluster1::> vserver services ldap client modify -client-config corp -vserver vs1 -bind-dn administrator@example.com
-user-dn "ou=People\;,dc=mypc,dc=example,dc=com; ou=People1,dc=mypc,dc=example2,dc=com"

```
The following example modifies an existing LDAP client configuration with
multiple user DNs, one of them containing a semicolon and a backslash.
```

cluster1::> vserver services ldap client modify -client-config corp -vserver vs1 -bind-dn administrator@example.com
-user-dn "ou=People\;,dc=mypc,dc=example,dc=com\\; ou=People1,dc=mypc,dc=example2,dc=com"

```
The following example modifies an existing LDAP client configuration with
netgroup by host DN.
```

cluster1::*>vserver services ldap client modify -vserver vs1 -client-config corp

-netgroup-byhost-dn nisMapName="netgroup.byhost",dc=rfcbis,dc=com

```
== Related Links

* xref:{relative_path}vserver-services-name-service-ldap-client-modify-
bind-password.html[vserver services name-service ldap client modify-bind-
password]




[[IDfb75b795fead4c78403c8714091cb081]]
= vserver services name-service ldap client show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display LDAP client configurations

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap client show` command displays
information about LDAP client configurations which a Vserver can be
associated with. An LDAP client configuration created by a Vserver's
administrator or by the cluster administrator for the Vserver is owned by
the Vserver. A cluster-wide LDAP client configuration is created by a
cluster administrator by specifying the admin Vserver's name as a value to
the -vserver parameter. In addition to its owned LDAP client
configurations, a Vserver can be associated with such cluster-wide LDAP
client configurations.

== 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.
```

```
``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays all LDAP client
configurations that can be associated with the specified Vserver. A data
Vserver or admin Vserver can be specified.

``[-client-config <text>]`` - Client Configuration Name::
If you specify this parameter, the command displays information about the
LDAP client configuration you specify.

``[-ldap-servers <text>,...]`` - LDAP Server List::
If you specify this parameter, the command displays LDAP client
configurations using the specified list of LDAP servers.

``[-servers <IP Address>,...]`` - (DEPRECATED)-LDAP Server List::
(DEPRECATED)-If you specify this parameter, the command displays LDAP
client configurations using the specified list of LDAP servers.

``[-ad-domain <TextNoCase>]`` - Active Directory Domain::
If you specify this parameter, the command displays LDAP client
configurations using the specified domain to discover their list of LDAP
servers.

``[-preferred-ad-servers <IP Address>,...]`` - Preferred Active Directory
Servers::
If you specify this parameter, the command displays LDAP client
configurations using the specified list of preferred servers.

``[-restrict-discovery-to-site {true|false}]`` - Restrict discovery to
site scope::
If you specify this parameter, the command displays only the LDAP client
configurations that do site-scope discovery.

``[-bind-as-cifs-server {true|false}]`` - Bind Using the Vserver's CIFS
Credentials::
If you specify this parameter, the command displays LDAP client
configurations that bind using CIFS server credentials. If the CIFS server
is in workgroup mode, the value of this parameter should be false.

``[-schema <text>]`` - Schema Template::
If you specify this parameter, the command displays LDAP client
configurations using the specified schema.

``[-port <integer>]`` - LDAP Server Port::
If you specify this parameter, the command displays LDAP client
configurations using the specified server port.
```

```
``[-query-timeout <integer>]`` - Query Timeout (sec)::
If you specify this parameter, the command displays LDAP client
configurations using the specified query timeout (in seconds).

``[-min-bind-level {anonymous|simple|sasl}]`` - Minimum Bind
Authentication Level::
If you specify this parameter, the command displays LDAP client
configurations using the specified minimum bind level.

``[-bind-dn <ldap_dn>]`` - Bind DN (User)::
If you specify this parameter, the command displays LDAP client
configurations using the specified bind DN.

``[-base-dn <ldap_dn>]`` - Base DN::
If you specify this parameter, the command displays LDAP client
configurations using the specified base DN.

``[-base-scope {base|onelevel|subtree}]`` - Base Search Scope::
If you specify this parameter, the command displays LDAP client
configurations using the specified base search scope.

``[-user-dn <ldap_dn>]`` - User DN [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified user DN.

``[-user-scope {base|onelevel|subtree}]`` - User Search Scope [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified user search scope.

``[-group-dn <ldap_dn>]`` - Group DN [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified group DN.

``[-group-scope {base|onelevel|subtree}]`` - Group Search Scope [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified group search scope.

``[-netgroup-dn <ldap_dn>]`` - Netgroup DN [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified netgroup DN.

``[-netgroup-scope {base|onelevel|subtree}]`` - Netgroup Search Scope
```

```
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified netgroup search scope.

``[-is-owner {true|false}]`` - Vserver Owns Configuration::
If you set this parameter to true, the command displays LDAP client
configurations with the Vservers which own them.

``[-use-start-tls {true|false}]`` - Use start-tls Over LDAP Connections::
This parameter specifies whether or not to use Start TLS over LDAP
connections. When enabled, the communication between the ONTAP LDAP Client
and the LDAP Server will be encrypted using Start TLS. Start TLS is a
mechanism to provide secure communication by using the TLS/SSL protocols.
If you do not specify this parameter, the default is `false` .

``[-is-netgroup-byhost-enabled {true|false}]`` - Enable Netgroup-By-Host
Lookup [.advaced-priv]#(privilege: advanced)#::
If you set this parameter to true, the command displays LDAP client
configurations for which netgroup-by-host lookup is enabled.

``[-netgroup-byhost-dn <ldap_dn>]`` - Netgroup-By-Host DN [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified netgroup-by-host DN.

``[-netgroup-byhost-scope {base|onelevel|subtree}]`` - Netgroup-By-Host
Scope [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified netgroup-by-host search scope.

``[-session-security {none|sign|seal}]`` - Client Session Security::
If this parameter is set to seal, the command displays LDAP client
configurations where both signing and sealing are required for LDAP
communications. If set to sign, the command displays LDAP client
configurations where only signing is required for LDAP communications. If
set to none, the command displays LDAP client configurations where no
security is required for LDAP communications.

``[-referral-enabled {true|false}]`` - LDAP Referral Chasing::
If you specify this parameter, the command displays information about LDAP
referral configurations using the specified client.

``[-group-membership-filter <text>]`` - Group Membership Filter [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays LDAP client
configurations using the specified group-membership filter.
```

```
``[-ldaps-enabled {true|false}]`` - Is LDAPS Enabled::
If you specify this parameter, the command displays LDAP client
configurations using the specified value of this parameter.

``[-try-channel-binding {true|false}]`` - Try Channel Binding::
If you specify this parameter, the command displays LDAP client
configurations using the specified channel binding.

== Examples

The following example shows a summary of all of the LDAP client
configurations available for Vserver `vs1` :
```

cluster1::> vserver services name-service ldap client show -vserver vs1
Vserver Client LDAP Active Directory Minimum
Configuration Servers Domain Schema Bind Level
---------- ---------- ---------- ---------- ---------- ----------
vs1 corp ldapserver. - RFC-2307 anonymous
example.com
vs1 corpnew 172.16.0.200 - RFC-2307 simple

```
[[ID9d66bb7fab53e628b25b7621fccafc9a]]
= vserver services name-service ldap client schema copy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Copy an existing LDAP schema template

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service ldap client schema copy` command
creates a new LDAP schema template from an existing one. In addition to an
owned LDAP schema template, a Vserver administrator can also copy a
cluster-wide LDAP schema template that is owned by the admin Vserver.

== Parameters

``[-vserver <Vserver Name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver for which you want to copy an
existing LDAP schema template.

``-schema <text>`` - Schema Template [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the existing schema template you want
to copy.

``-new-schema-name <text>`` - New Schema Template Name [.advaced-
priv]#(privilege: advanced)#::
This parameter specifies the name of the schema template copy.

== Examples

The following example creates a copy of the `RFC-2307` schema template and
names it `corp-schema` for Vserver "vs1":
```

cluster1::> vserver services name-service ldap client schema copy -vserver vs1 -schema RFC-2307 -new -schema-name corp-schema

```
[[IDe0cf8955a51681675549ed3e605eee5a]]
= vserver services name-service ldap client schema delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an LDAP schema template

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service ldap client schema delete` command
deletes an LDAP schema template. A Vserver administrator can only delete
templates owned by the Vserver.

NOTE: You cannot delete the default schema templates.



== Parameters

``[-vserver <Vserver Name>]`` - Vserver::
This parameter specifies the name of Vserver owning the LDAP schema
template you want to delete.

``-schema <text>`` - Schema Template::
This parameter specifies the name of the schema template you want to
delete.

== Examples

The following example deletes a schema template named `corp-schema` owned
by Vserver `vs1` :
```

cluster1::> vserver services name-service ldap client schema delete -vserver vs1 -schema corp-schema

```
[[ID0d669019cd2dbd14b8a9b428755acc5b]]
= vserver services name-service ldap client schema modify
:hardbreaks:
:linkattrs:
:relative_path: ./
```

```
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an LDAP schema template

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service ldap client schema modify` command
modifies an existing LDAP schema template. You cannot modify the default
schema templates. Create a copy of a default schema template using the
xref:{relative_path}vserver-services-name-service-ldap-client-schema-
copy.html[vserver services name-service ldap client schema copy] command,
and then modify the copy. A Vserver administrator can only modify
templates owned by the Vserver.

== Parameters

``[-vserver <Vserver Name>]`` - Vserver::
This parameter specifies the name of the Vserver owning the LDAP schema
template you want to modify.

``-schema <text>`` - Schema Template::
This parameter specifies the name of the schema template you want to
modify.

``[-comment <text>]`` - Comment::
This parameter specifies a comment that describes the schema template.

``[-posix-account-object-class <text>]`` - RFC 2307 posixAccount Object
Class::
This parameter specifies the RFC 2307 posixAccount object class name
defined by the schema.

``[-posix-group-object-class <text>]`` - RFC 2307 posixGroup Object
Class::
This parameter specifies the RFC 2307 posixGroup object class name defined
by the schema.

``[-nis-netgroup-object-class <text>]`` - RFC 2307 nisNetgroup Object
Class::
This parameter specifies the RFC 2307 nisNetgroup object class name
defined by the schema.

``[-uid-attribute <text>]`` - RFC 2307 uid Attribute::
```

```
This parameter specifies the RFC 2307 uid attribute name defined by the
schema.

``[-uid-number-attribute <text>]`` - RFC 2307 uidNumber Attribute::
This parameter specifies the RFC 2307 uidNumber attribute name defined by
the schema.

``[-gid-number-attribute <text>]`` - RFC 2307 gidNumber Attribute::
This parameter specifies the RFC 2307 gidNumber attribute name defined by
the schema.

``[-cn-group-attribute <text>]`` - RFC 2307 cn (for Groups) Attribute::
This parameter specifies the RFC 2307 cn (for Groups) attribute name
defined by the schema.

``[-cn-netgroup-attribute <text>]`` - RFC 2307 cn (for Netgroups)
Attribute::
This parameter specifies the RFC 2307 cn (for Netgroups) attribute name
defined by the schema.

``[-user-password-attribute <text>]`` - RFC 2307 userPassword Attribute::
This parameter specifies the RFC 2307 userPassword attribute name defined
by the schema.

``[-gecos-attribute <text>]`` - RFC 2307 gecos Attribute::
This parameter specifies the RFC 2307 gecos attribute name defined by the
schema.

``[-home-directory-attribute <text>]`` - RFC 2307 homeDirectory
Attribute::
This parameter specifies the RFC 2307 homeDirectory attribute name defined
by the schema.

``[-login-shell-attribute <text>]`` - RFC 2307 loginShell Attribute::
This parameter specifies the RFC 2307 loginShell attribute name defined by
the schema.

``[-member-uid-attribute <text>]`` - RFC 2307 memberUid Attribute::
This parameter specifies the RFC 2307 memberUid attribute name defined by
the schema.

``[-member-nis-netgroup-attribute <text>]`` - RFC 2307 memberNisNetgroup
Attribute::
This parameter specifies the RFC 2307 memberNisNetgroup attribute name
defined by the schema.

``[-nis-netgroup-triple-attribute <text>]`` - RFC 2307 nisNetgroupTriple
```

```
Attribute::
This parameter specifies the RFC 2307 nisNetgroupTriple attribute name
defined by the schema.

``[-enable-rfc2307bis {true|false}]`` - Enable Support for Draft RFC
2307bis::
This parameter specifies whether RFC 2307bis is enabled for the schema.

``[-group-of-unique-names-object-class <text>]`` - RFC 2307bis
groupOfUniqueNames Object Class::
This parameter specifies the RFC 2307bis groupOfUniqueNames object class
name defined by the schema. This parameter takes effect only when RFC
2307bis is enabled for the schema.

``[-unique-member-attribute <text>]`` - RFC 2307bis uniqueMember
Attribute::
This parameter specifies the RFC 2307bis uniqueMember attribute name
defined by the schema. This parameter takes effect only when RFC 2307bis
is enabled for the schema.

``[-windows-to-unix-object-class <text>]`` - Data ONTAP Name Mapping
windowsToUnix Object Class::
This parameter specifies the name mapping windowsToUnix object class name
defined by the schema.

``[-windows-account-attribute <text>]`` - Data ONTAP Name Mapping
windowsAccount Attribute::
This parameter specifies the name mapping windowsAccount attribute name
defined by the schema.

``[-windows-to-unix-attribute <text>]`` - Data ONTAP Name Mapping
windowsToUnix Attribute::
This parameter specifies the name mapping windowsToUnix attribute name
defined by the schema.

``[-windows-to-unix-no-domain-prefix {true|false}]`` - No Domain Prefix
for windowsToUnix Name Mapping::
This parameter specifies the name mapping windowsToUnixNoDomainPrefix
setting defined by the schema.

``[-nis-object-class <text>]`` - RFC 2307 nisObject Object Class::
This parameter specifies the nisObject class name defined by the schema.
This parameter takes effect only when netgroup.byhost is enabled for the
vserver.

``[-nis-mapname-attribute <text>]`` - RFC 2307 nisMapName Attribute::
```

```
This parameter specifies the nisMapName attribute name defined by the
schema. This parameter takes effect only when netgroup.byhost is enabled
for the vserver.

``[-nis-mapentry-attribute <text>]`` - RFC 2307 nisMapEntry Attribute::
This parameter specifies the nisMapEntry attribute name defined by the
schema. This parameter takes effect only when netgroup.byhost is enabled
for the vserver.

== Examples

The following example modifies the schema template called `corp-schema`
owned by Vserver `vs1` to use `User` as the uid attribute name:
```

cluster1::> vserver services name-service ldap client schema modify -vserver vs1 -schema corp-schema -uid -attribute User

```
== Related Links

* xref:{relative_path}vserver-services-name-service-ldap-client-schema-
copy.html[vserver services name-service ldap client schema copy]




[[ID08e788bcd8b8c26b7f5b2e7dc676c751]]
= vserver services name-service ldap client schema show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display LDAP schema templates

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ldap client schema show` command shows
information about LDAP schema templates which a Vserver can access. An
LDAP schema template created by a Vserver's administrator or by the
cluster administrator for the Vserver is owned by the Vserver. A cluster-
wide LDAP schema template is created by a cluster administrator by
specifying the admin Vserver's name as a value to the -vserver parameter.
```

```
In addition to its owned LDAP schema templates, a Vserver can access such
cluster-wide LDAP schema templates.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
If you specify this parameter, the command displays all LDAP schema
templates that can be accessed by the specified Vserver.

``[-schema <text>]`` - Schema Template::
If you specify this parameter, the command displays the schema template
with the specified name.

``[-comment <text>]`` - Comment::
If you specify this parameter, the command displays schema templates with
the specified comment.

``[-posix-account-object-class <text>]`` - RFC 2307 posixAccount Object
Class::
If you specify this parameter, the command displays schema templates with
the specified posixAccount object class.

``[-posix-group-object-class <text>]`` - RFC 2307 posixGroup Object
Class::
If you specify this parameter, the command displays schema templates with
the specified posixGroup object class.

``[-nis-netgroup-object-class <text>]`` - RFC 2307 nisNetgroup Object
Class::
If you specify this parameter, the command displays schema templates with
the specified nisNetgroup object class.

``[-uid-attribute <text>]`` - RFC 2307 uid Attribute::
If you specify this parameter, the command displays schema templates with
the specified uid attribute.

``[-uid-number-attribute <text>]`` - RFC 2307 uidNumber Attribute::
If you specify this parameter, the command displays schema templates with
```

```
the specified uidNumber attribute.

``[-gid-number-attribute <text>]`` - RFC 2307 gidNumber Attribute::
If you specify this parameter, the command displays schema templates with
the specified gidNumber attribute.

``[-cn-group-attribute <text>]`` - RFC 2307 cn (for Groups) Attribute::
If you specify this parameter, the command displays schema templates with
the specified cn (for Groups) attribute.

``[-cn-netgroup-attribute <text>]`` - RFC 2307 cn (for Netgroups)
Attribute::
If you specify this parameter, the command displays schema templates with
the specified cn (for Netgroups) attribute.

``[-user-password-attribute <text>]`` - RFC 2307 userPassword Attribute::
If you specify this parameter, the command displays schema templates with
the specified userPassword attribute.

``[-gecos-attribute <text>]`` - RFC 2307 gecos Attribute::
If you specify this parameter, the command displays schema templates with
the specified gecos attribute.

``[-home-directory-attribute <text>]`` - RFC 2307 homeDirectory
Attribute::
If you specify this parameter, the command displays schema templates with
the specified homeDirectory attribute.

``[-login-shell-attribute <text>]`` - RFC 2307 loginShell Attribute::
If you specify this parameter, the command displays schema templates with
the specified loginShell attribute.

``[-member-uid-attribute <text>]`` - RFC 2307 memberUid Attribute::
If you specify this parameter, the command displays schema templates with
the specified memberUid attribute.

``[-member-nis-netgroup-attribute <text>]`` - RFC 2307 memberNisNetgroup
Attribute::
If you specify this parameter, the command displays schema templates with
the specified memberNisNetgroup attribute.

``[-nis-netgroup-triple-attribute <text>]`` - RFC 2307 nisNetgroupTriple
Attribute::
If you specify this parameter, the command displays schema templates with
the specified nisNetgroupTriple attribute.

``[-enable-rfc2307bis {true|false}]`` - Enable Support for Draft RFC
```

```
2307bis::
If you set this parameter to true, the command displays RFC 2307bis
enabled LDAP schema templates.

``[-group-of-unique-names-object-class <text>]`` - RFC 2307bis
groupOfUniqueNames Object Class::
If you specify this parameter, the command displays schema templates with
the specified groupOfUniqueNames object class.

``[-unique-member-attribute <text>]`` - RFC 2307bis uniqueMember
Attribute::
If you specify this parameter, the command displays schema templates with
the specified uniqueMember attribute.

``[-windows-to-unix-object-class <text>]`` - Data ONTAP Name Mapping
windowsToUnix Object Class::
If you specify this parameter, the command displays schema templates with
the specified windowsToUnix object class.

``[-windows-account-attribute <text>]`` - Data ONTAP Name Mapping
windowsAccount Attribute::
If you specify this parameter, the command displays schema templates with
the specified windowsAccount attribute.

``[-windows-to-unix-attribute <text>]`` - Data ONTAP Name Mapping
windowsToUnix Attribute::
If you specify this parameter, the command displays schema templates with
the specified windowsToUnix attribute.

``[-windows-to-unix-no-domain-prefix {true|false}]`` - No Domain Prefix
for windowsToUnix Name Mapping::
If you specify this parameter, the command displays schema templates with
the specified windowsToUnixNoDomainPrefix setting.

``[-is-owner {true|false}]`` - Vserver Owns Schema::
If you set this parameter to true, the command displays LDAP schema
templates with the Vservers which own them.

``[-nis-object-class <text>]`` - RFC 2307 nisObject Object Class::
If you specify this parameter, the command displays schema templates with
the specified nisObject attribute.

``[-nis-mapname-attribute <text>]`` - RFC 2307 nisMapName Attribute::
If you specify this parameter, the command displays schema templates with
the specified nisMapName attribute.

``[-nis-mapentry-attribute <text>]`` - RFC 2307 nisMapEntry Attribute::
```

```
If you specify this parameter, the command displays schema templates with
the specified nisMapEntry attribute.

== Examples

The following example shows a summary of all of the default LDAP schema
templates defined in the cluster:
```

cluster1::> vserver services name-service ldap client schema show
Vserver Schema Template Comment
------- --------------- --
cluster-node3
MS-AD-BIS Schema based on Active Directory Identity Management for UNIX (read-only)
cluster-node3
AD-IDMU Schema based on Active Directory Identity Management for UNIX (read-only)
cluster-node3
AD-SFU Schema based on Active Directory Services for UNIX (read-only)
cluster-node3
RFC-2307 Schema based on RFC 2307 (read-only)
4 entries were displayed.

```
[[ID4408f4b3d5b73e10d86ec58f98029892]]
= vserver services name-service netgroup load
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Load netgroup definitions from a URI

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service netgroup load` command loads netgroup
definitions from a uniform resource identifier (URI) to a specified
Vserver. You can load from a netgroup file at an FTP or a HTTP location
(source URI) using the respective protocol.

Before ONTAP saves the new netgroup definitions, it checks that the
netgroup file does not have any file structure issues, does not contain
any syntax errors, and all entries comply with the following rules:

* A domain name consists of one or more labels separated by periods (.).
* A hostname is a valid domain name, IPv4 address, or IPv6 address.
```

```
* Valid characters for a label are all alphanumeric characters, underscore
(_), and dash (-). A label may not begin or end with a dash.
* Valid characters for a username are all ASCII printable characters with
the exception of whitespace, parentheses, and comma (,).
* Valid characters for a netgroup name are all alphanumeric characters,
underscore (_), and dash (-). A netgroup name may not begin with a dash.
* A single line in the netgroup file may not exceed 4096 characters.



If the file is found to contain errors, ONTAP will issue an error to that
effect and netgroup definitions will not be loaded into the specified
Vserver. After correcting the error, reload the netgroup file into the
specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver for which you want to load netgroup
definitions.

``-source {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}`` - URI
to Load from::
This parameter specifies the source URI from which you want to load
netgroup definitions. You can load from a URI either using the FTP or the
HTTP protocol.

``[-foreground {true|false}]`` - Load Netgroup in the Foreground::
This parameter specifies whether the operation runs in the foreground. The
default setting is `true` (the operation runs in foreground). When set to
true, the command does not return until the operation completes.

``[-skip-hostname-validation <true>]`` - Skip Hostname Validation
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the hostname validation is skipped.

``[-skip-file-size-check <true>]`` - Skip File Size Check Before Download
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the file is downloaded without checking
the file size. Use this parameter if the server does not supply the file
size or does not provide an accurate value. This parameter can also be
used to download a file greater than the default 5 MB size limit.
+
NOTE: If this parameter is specified and the file is very large, the
transfer may take a long time or fail due to disk space limitations.
+
```

```
``[-skip-file-duplicate-check <true>]`` - Skip Netgroup File Duplicate
Check [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the netgroup file is downloaded even if
the contents are same as the existing netgroup file. In this case, the
existing file will be replaced.

== Examples

The following example loads netgroup definitions into a Vserver named vs1
from the file netgroup1 at FTP location ftp://ftp.example.com:
```

cluster1::> vserver services name-service netgroup load -vserver vs1 -source ftp://ftp.example.com/netgroup1

```
[[IDcd44b8b8634ce8fca58204e6a3165ae4]]
= vserver services name-service netgroup status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local netgroup definitions status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service netgroup status` command displays the
status of local netgroup definitions across a cluster. This enables you to
verify that netgroup definitions are consistent across all nodes that back
a Vserver into which netgroup definitions have been loaded.

The command displays the following information:

* Vserver name
* Node name
* Load time for netgroup definitions
* Hash value of the netgroup definitions
* Hash value of the netgroup-by-host database
* File size of the netgroup definitions file



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the netgroup status
only for the specified Vserver.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the netgroup status
only for the specified node.

``[-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Load Time [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the netgroup definitions that were loaded at the specified time. Specify
time in the format MM/DD/YYYY HH:MM:SS. Note that the load time stamps for
identical definitions are different on different nodes, because each node
downloads the definitions from the URI individually.

``[-hashvalue <text>]`` - Hash Value [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, command displays the status only for the
netgroup definitions that have the specified hash value. Note that the
primary purpose of the command is to verify that the definitions on all
nodes have the same hash value, so querying on a specific hash value is
not useful in most cases.

``[-hashvalue-byhost <text>]`` - Hash Value Byhost [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the netgroup definitions that have the specified hash value for netgroup-
by-host database. Note that the primary purpose of the command is to
verify that the definitions on all nodes have the same hash value for
netgroup-by-host database.

``[-filesize {<integer>[KB|MB|GB|TB|PB]}]`` - File Size [.advaced-
priv]#(privilege: advanced)#::
```

```
If you specify this parameter, the command displays the status only for
the netgroup definitions that have the specified file size. Note that the
primary purpose of the command is to verify that the definitions on all
nodes have the same file size, so querying on a specific file size is not
useful in most cases.

== Examples

The following example displays the netgroup definition status for all
Vservers:
```

cluster1::*> vserver services name-service netgroup status
Vserver Node Load Time Hash Value Hash Value By-Host File Size
--------- ------- ------------------- ------------------------------- ------------------------------- ----------
vs1
node1 9/20/2008 16:04:55 e6cb38ec1396a280c0d2b77e3a84eda2 913a182a72aa1872495be398ebb2cd23
1.00KB
node2 9/20/2008 16:04:53 e6cb38ec1396a280c0d2b77e3a84eda2 913a182a72aa1872495be398ebb2cd23
1.00KB
vs2
node1 9/20/2008 16:06:26 c0d2b77e3a84eda2e6cb38ec1396a280 009321eddb45611e95d9f7f277ec0621
2.3MB
node2 9/20/2008 16:06:27 c0d2b77e3a84eda2e6cb38ec1396a280 009321eddb45611e95d9f7f277ec0621
2.3MB
4 entries were displayed.

```
[[ID235b279d9b3c1a7551de96ca0b993fa8]]
= vserver services name-service netgroup file delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a local netgroup file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service netgroup file delete` command deletes
the local netgroup files for given Vservers.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vservers whose local netgroup file you
want to delete. Separate multiple Vserver names with commas.

== Examples

The following example deletes the local netgroup file for a Vserver named
vs1.
```

cluster1::> vserver services netgroup file delete -vserver vs1

```
[[IDced2e066a3edc415c35a7b9eba8cec41]]
= vserver services name-service netgroup file show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display a local netgroup file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
The `vserver services netgroup file show` command displays the contents of
the local netgroup file for the specified Vservers. All the entries under
a given netgroup, specified in the Netgroup column of the command output,
list the members of that netgroup. Each netgroup file specifies netgroups,
which are sets of tuples. Each member of a netgroup is either the name of
another netgroup, specified in the Member Netgroup column, or a
specification of a tuple as follows: (Host, User, Domain) where Host,
User, and Domain are character string names for the corresponding
component. Any of the components of a tuple can either be empty to specify
a wildcard value or a dash (-) to specify no valid value.

== 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.

``[-vserver <vserver name>]`` - Vserver::
Use this parameter to display the local netgroup file contents for the
Vservers you specify.

``[-netgroup <text>]`` - Netgroup Name::
If you specify this parameter, the command displays information about the
netgroup you specify.

``[-netgrpmemb <text>]`` - Member Netgroup::
If you specify this parameter, the command displays information about the
member netgroup you specify.

``[-host <text>]`` - Member Host::
If you specify this parameter, the command displays information about the
host you specify.

``[-user <text>]`` - Member User::
If you specify this parameter, the command displays information about the
user you specify.

``[-domain <text>]`` - Member Domain::
If you specify this parameter, the command displays information about the
domain you specify.
```

```
== Examples

The following example displays the netgroup file contents for the Vserver
named vs1.
```

cluster1::> vserver services netgroup file show -vserver vs1
Member
Vserver Netgroup Netgroup Host User Domain
---------- -------- ---------- ---------- ---------- ----------
vs1
netgrp1
netgrp9 - - -
h1 d1
h22 d22
netgrp11 - - -
netgrp18 - - -
h119 u4343 d34
netgrp8
'-' u88 '-'

```
[[IDba00467049070179754487 4fa966fdf6]]
= vserver services name-service nis-domain create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a NIS domain configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service nis-domain create` command creates a
configuration for an NIS domain. You can configure only one NIS domain for
a given Vserver. You can also configure more than one Vserver with the
same NIS domain.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Use this parameter to specify the Vserver on which the NIS domain
configuration is created. A data Vserver or admin Vserver can be
specified.
```

```
``-domain <nis domain>`` - NIS Domain::
Use this parameter to specify the NIS domain for which a configuration is
created. Maximum Supported NIS Domain length: 64 characters.

{ ``-nis-servers <text>,...`` - NIS Servers::
Use this parameter to specify the hostnames/IP addresses of NIS servers
used by the NIS domain configuration. Separate multiple hostnames/IP
addresses with commas.

| ``-servers <IP Address>,...`` - (DEPRECATED)-NIS Server }::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP.
+
Use this parameter to specify the IP addresses of NIS servers used by the
NIS domain configuration. Separate multiple IP addresses with commas.

== Examples

The following example creates an NIS domain configuration on the Vserver
named vs0. The NIS domain is named nisdomain and uses an NIS server with
the IP address 192.0.2.180.
```

cluster1::> vserver services name-service nis-domain create -vserver vs0 -domain nisdomain -nis-servers 192.0.2.180

```
[[IDa5532f5d17e3d1642da5ca824f76d5a7]]
= vserver services name-service nis-domain delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a NIS domain configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service nis-domain delete` command deletes an
NIS domain configuration.

Deleting a NIS domain configuration removes it permanently.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Use this parameter to specify the Vserver from which the NIS domain
configuration is deleted. A data Vserver or admin Vserver can be
specified.

``-domain <nis domain>`` - NIS Domain::
Use this parameter to specify the NIS domain whose configuration is
deleted.

== Examples

The following example deletes the configuration of an NIS domain named
testnisdomain from a Vserver named vs2:
```

cluster1::> vserver services name-service nis-domain delete -vserver vs2 -domain testnisdomain

```
[[IDc30666bfc0bfbb8d781f22802b938e9d]]
= vserver services name-service nis-domain modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Modify a NIS domain configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service nis-domain modify` command to
modify the NIS server of a NIS domain configuration.

To change the NIS domain, delete the NIS configuration using the
xref:{relative_path}vserver-services-name-service-nis-domain-
delete.html[vserver services name-service nis-domain delete] command and
then create the NIS configuration with new NIS domain using the
xref:{relative_path}vserver-services-name-service-nis-domain-
create.html[vserver services name-service nis-domain create] command. To
permanently remove a configuration, use the xref:{relative_path}vserver-
services-name-service-nis-domain-delete.html[vserver services name-service
nis-domain delete] command.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Use this parameter to specify the Vserver whose NIS domain configuration
is modified. A data Vserver or admin Vserver can be specified.

``-domain <nis domain>`` - NIS Domain::
Use this parameter to specify the NIS domain whose configuration is
modified.

{ ``[-nis-servers <text>,...]`` - NIS Servers::
Use this parameter to specify the hostnames/IP addresses of NIS servers
used by the the NIS domain configuration. Separate multiple hostnames/IP
addresses with commas.

| ``[-servers <IP Address>,...]`` - (DEPRECATED)-NIS Server }::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP.
+
Use this parameter to specify the IP addresses of NIS servers used by the
the NIS domain configuration. Separate multiple IP addresses with commas.

== Examples

The following example modifies the NIS servers of a NIS domain named
```

```
nisdomain on a Vserver named vs0:
```

cluster1::> vserver services name-service nis-domain modify -vserver vs0 -domain nisdomain -nis-servers 192.0.2.180

```
== Related Links

* xref:{relative_path}vserver-services-name-service-nis-domain-
delete.html[vserver services name-service nis-domain delete]
* xref:{relative_path}vserver-services-name-service-nis-domain-
create.html[vserver services name-service nis-domain create]




[[IDf0a604aa82d08d98411aa7ced9cebdd5]]
= vserver services name-service nis-domain show-bound
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display binding status of a NIS domain configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service nis-domain show-bound` command displays
binding information about NIS domain configurations.

== 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.

``[-vserver <vserver name>]`` - Vserver::
If you use this parameter, the command displays binding information only
```

```
about the NIS domain configurations of the specified Vservers. Use this
parameter with the `-domain` parameter to display binding information only
about a particular NIS domain configuration on the specified Vserver. A
data Vserver or admin Vserver can be specified.

``[-domain <nis domain>]`` - NIS Domain::
If you use this parameter, the command displays binding information only
about the NIS domain configurations that match the specified NIS domain
name. Use this parameter with the `-vserver` parameter to display binding
information only about a particular Vserver on the specified NIS domain
name.

``[-bound-servers <IP Address>,...]`` - Bound NIS Servers::
If you use this parameter, the command displays NIS binding information
only about the specified NIS servers.

== Examples

The following example displays binding information about all NIS domain
configurations:
```

cluster1::> vserver services name-service nis-domain show-bound
Bound
Vserver Domain NIS Server
---------- ---------- ----------
vs1 testnisdomain1 192.0.2.180,
10.0.2.15
vs2 testnisdomain2 10.0.2.17
2 entries were displayed.

```
[[ID9a871bd2d9e1d55324f93be8e31b060e]]
= vserver services name-service nis-domain show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NIS domain configurations

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service nis-domain show` command displays
information about NIS domain configurations.
```

```
== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Use this parameter to display information only about the NIS domain
configurations of the Vservers you specify. Use this parameter with the `-
domain` parameter to display information only about a particular NIS
domain configuration on the Vserver you specify. A data Vserver or admin
Vserver can be specified.

``[-domain <nis domain>]`` - NIS Domain::
Use this parameter to display information only about the NIS domain
configurations that match the NIS domain name you specify. Use this
parameter with the `-vserver` parameter to display information only about
a particular NIS domain configuration on the Vserver you specify.

``[-nis-servers <text>,...]`` - NIS Servers::
Use this parameter to display information only about the NIS domain
configurations that use the NIS servers at the hostnames/IP addresses you
specify.

``[-servers <IP Address>,...]`` - (DEPRECATED)-NIS Server::
+
NOTE: This parameter has been deprecated and might be removed in a future
version of ONTAP.
+
Use this parameter to display information only about the NIS domain
configurations that use the NIS servers at the IP addresses you specify.

== Examples

The following example displays information about all NIS domain
configurations:
```

cluster1::> vserver services name-service nis-domain show
Vserver Domain NIS Server
---------- ---------- ----------

vs1 nisdomain 192.0.2.180
vs2 nisdomain 10.0.2.15
vs3 testnisdomain 192.0.2.128, 192.0.2.180
3 entries were displayed.

```
[[IDec301e689a971abce18c2ded3e3ef8ce]]
= vserver services name-service nis-domain group-database build
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Build NIS group database

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service nis-domain group-database build`
command rebuilds the NIS group.byuser DB for a given Vserver if NIS is
added as source for group and an active nis-domain exists.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which NIS group.byuser DB
will be rebuilt. A data Vserver can be specified.

== Examples

The following example rebuilds NIS group.byuser DB for Vserver vs0.
```

cluster1::> vserver services name-service nis-domain group-database build -vserver vs0

```
[[ID82a123eaa2be889b1f10eb76996e6caf]]
= vserver services name-service nis-domain group-database status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NIS group database status of the local node
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service nis-domain group-database status`
command displays the status of local NIS group.byuser db across a cluster.
This enables you to verify that NIS group.byuser db are consistent across
all nodes.

The command displays the following information:

* Vserver name
* Node name
* Last build time of NIS group.byuser db
* Hash value of the NIS group.byuser db
* File size of the NIS group.byuser db



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the NIS group.byuser
db status only for the specified Vserver.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the NIS group.byuser
db status only for the specified node.

``[-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Load Time [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the NIS group.byuser db that were built at the specified time. Specify
time in the format MM/DD/YYYY HH:MM:SS. Note that the load time stamps for
```

```
identical definitions are different on different nodes, because each node
extracts the db individually.

``[-filesize {<integer>[KB|MB|GB|TB|PB]}]`` - File Size [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the NIS group.byuser db that have the specified file size. Note that the
primary purpose of the command is to verify that the definitions on all
nodes have the same file size, so querying on a specific file size is not
useful in most cases.

``[-hashvalue <text>]`` - Hash Value [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, command displays the status only for the
NIS group.byuser db that have the specified hash value. Note that the
primary purpose of the command is to verify that the definitions on all
nodes have the same hash value, so querying on a specific hash value is
not useful in most cases.

== Examples

The following example displays the NIS group.byuser db status for vserver
vs0 :
```

cluster1::*> vserver services name-service nis-domain group-database status -vserver vs0
Vserver Node Last Build Time File Size
--------- --------------- ------------------ ----------
Hash Value
----
vs0
node1 2/14/2017 11:39:56 136KB
a30b7d6d03197a7af25de72dcc4bd64f

```
[[IDac47b46a476beb151d84e75ef4cef253]]
= vserver services name-service nis-domain netgroup-database build
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Build NIS netgroup database

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service nis-domain netgroup-database build`
command rebuilds the NIS netgroup.byhost database for a given Vserver if
NIS is added as the source for netgroup and an active NIS domain exists.

== Parameters

``-vserver <vserver name>`` - Vserver Name [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the NIS
netgroup.byhost database will be rebuilt. A data Vserver can be specified.

== Examples

The following example rebuilds the NIS netgroup.byhost database for
Vserver vs0.
```

cluster1::> vserver services name-service nis-domain netgroup-database build -vserver vs0

```
[[ID0b221955f6b0c5be6411032431e38c9c]]
= vserver services name-service nis-domain netgroup-database show-status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NIS netgroup database status of the local node

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.
```

```
== Description

The `vserver services name-service nis-domain netgroup-database show-
status` command displays the status of local NIS netgroup.byhost databases
across a cluster. This enables you to verify that NIS netgroup.byhost
databases are consistent across all nodes.

The command displays the following information:

* Vserver name
* Node name
* Last build time of the NIS netgroup.byhost database
* Hash value of the NIS netgroup.byhost database
* File size of the NIS netgroup.byhost database



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the NIS
netgroup.byhost database status only for the specified Vserver.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the NIS
netgroup.byhost database status only for the specified node.

``[-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Load Time [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the NIS netgroup.byhost databases that were built at the specified time.
Specify time in the format MM/DD/YYYY HH:MM:SS. Note that the load time
stamps for identical definitions are different on different nodes, because
each node extracts the databases individually.
```

```
``[-filesize {<integer>[KB|MB|GB|TB|PB]}]`` - File Size [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the NIS netgroup.byhost databases that have the specified file size. Note
that the primary purpose of the command is to verify that the definitions
on all nodes have the same file size, so querying on a specific file size
is not useful in most cases.

``[-hashvalue <text>]`` - Hash Value [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the status only for
the NIS netgroup.byhost databases that have the specified hash value. Note
that the primary purpose of the command is to verify that the definitions
on all nodes have the same hash value, so querying on a specific hash
value is not useful in most cases.

== Examples

The following example displays the NIS netgroup.byhost database status for
vserver vs0 :
```

cluster1::*> vserver services name-service nis-domain netgroup-database status -vserver vs0
Vserver Node Last Build Time File Size
--------- -------------- ------------------ ----------
Hash Value
---
vs0
node1 2/14/2019 11:39:56 136KB
a30b7d6d03197a7af25de72dcc4bd64f

```
[[ID4269aaec1e4888ee30508b24d7490d44]]
= vserver services name-service nis-domain netgroup-database config modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify NIS netgroup database configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service nis-domain netgroup-database config
modify` command modifies NIS netgroup.byhost database configuration of the
specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to specify the Vserver for which the settings need to
be modified.

``[-state {enabled|disabled}]`` - State [.advaced-priv]#(privilege:
advanced)#::
Use this parameter to enable and disable NIS netgroup.byhost databases.
Default value is disabled.

``[-build-interval <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Build Interval [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the NIS netgroup.byhost database build time
interval. Default value is 24 hours and minimum value can be set to 5
minutes. Setting build interval below 1 hour may impact the performance.

== Examples

The following example enables NIS netgroup.byhost database and sets the
build interval to 24 hours for Vserver vs0:
```

cluster1::> vserver services name-service nis-domain netgroup-database config modify -vserver vs0 -build -interval 24h -state enabled

```
[[ID58e28c4a9f9cc453ba700ce672e2b049]]
= vserver services name-service nis-domain netgroup-database config show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NIS netgroup database configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service nis-domain netgroup-database config
show` command displays NIS netgroup.byhost database settings.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about NIS netgroup.byhost databases belonging to the specified Vservers.

``[-state {enabled|disabled}]`` - State [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays information only
about NIS netgroup.byhost databases with the specified state.

``[-build-interval <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Build Interval [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays information only
about NIS netgroup.byhost databases with the specified build time
interval.

== Examples
```

```
The following example displays the settings for Vserver vs0:
```

cluster1::> vserver services name-service nis-domain netgroup-database config show -vserver vs0
Vserver: vs0
Build Interval: 30m
State: disabled

```
[[ID7c185e195d757f7b77ea1636e8f64a7d]]
= vserver services name-service ns-switch create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a new Name Service Switch table entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service ns-switch create` command specifies the
order in which to lookup the name service sources, for a given Vserver and
name service database. Each name service database contains some
information regarding hosts, group, password, netgroup or namemap. Such a
database comes from one or more name service sources such as files, DNS,
LDAP or NIS.

Note: The ` vserver services name-service ns-switch ` command provides the
functionality of the /etc/nsswitch.conf file on UNIX systems. For more
information, see the UNIX man page for nsswitch.conf(5).

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver on which to create the new name
service switch entry

``-database {hosts|group|passwd|netgroup|namemap}`` - Name Service Switch
Database::
Name Service Switch Database Use this parameter to specify the name
service database for which the order of the source lookup is being
specified. This parameter can have the following values:
+
  *  hosts
```

```
  *  group
  *  passwd
  *  netgroup
  *  namemap



``-sources {files|dns|ldap|nis}`` - Name Service Source Order::
Name Service Source Order Use this parameter to specify the name service
sources and the order in which to look them up for the specified Vserver
and name service database. Each name service source in the list for this
parameter must be one of the following:
+
  *  files
  *  dns
  *  ldap
  *  nis



+
Separate multiple name service sources with commas.
+
For each database specified with the `-database` parameter, one or more
sources must be specified. The valid sources for each database type are
shown in the following table:
```

```
+----------+-------------------+
| Database |  Valid Sources    |
|----------+-------------------|
| hosts    |  files, dns       |
        | group    |  files, nis, ldap |
| passwd   |  files, nis, ldap |
        | netgroup |  files, nis, ldap |
| namemap  |  files, ldap      |
+----------+-------------------+
```

```
+
NOTE: If "files" is not specified as the default source for the "passwd"
or "group" database, ensure that default user and group entries for the
'passwd' and 'group' respectively are present in the source configured.
Default entries for "passwd" database: nobody, pcuser, root, sshd, toor,
daemon, operator, bin, tty, kmem, games, news, man, smmsp, mailnull, bind,
proxy, uucp, pop, www, admin, diag, autosupport. Default entries for
"group" database: wheel, daemon, kmem, sys, tty, operator, mail, bin,
news, man, games, ftp, staff, sshd, smmsp, mailnull, guest, bind, proxy,
authpf, _pflogd, _dhcp, uucp, dialer, network, audit, www, antivirus,
nogroup, nobody.
+


== Examples

The following example creates name service source ordering for the hosts
database on a Vserver named vs0. The order of looking up the sources is
specified as files followed by DNS.
```

cluster1::> vserver services name-service ns-switch create -vserver vs0 -database hosts -sources files,dns

```
The following example creates the name service source ordering for the
passwd database on a Vserver named vs1. The order of looking up the
sources is specified as files, NIS and LDAP.
```

cluster1::> vserver services nameservice ns-switch create -vserver vs1 -database passwd -sources files,nis,ldap

```
[[ID191e43be766b0a1e6af5a81fe90f7995]]
= vserver services name-service ns-switch delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove a Name Service Switch table entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the ` vserver services name-service ns-switch delete ` command to
permanently remove an existing name service switch entry.

== Parameters

``-vserver <vserver name>`` - Vserver::
Vserver Use this parameter to specify the Vserver for which to delete the
name service switch entry.

``-database {hosts|group|passwd|netgroup|namemap}`` - Name Service Switch
Database::
Name Service Switch Database Use this parameter to specify the name
service database, of the Vserver, for which the name service switch entry
is to be deleted. Following are the possible values for this parameter:
+
  *  hosts
  *  group
  *  passwd
  *  netgroup
  *  name_map



== Examples

The following example deletes the name service switch entry for the hosts
database on a Vserver named vs0.
```

cluster1::> vserver services name-service ns-switch delete -vserver vs0 -database hosts.

```
The following example deletes the name service switch entry for the group
database on a Vserver named vs1.
```

cluster1::> vserver services name-service ns-switch delete -vserver vs1 -database group.

```
[[ID0accce70dd69631e70b2f3055ebfea88]]
= vserver services name-service ns-switch modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change a Name Service Switch table entry

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

 Use the ` vserver services name-service ns-switch modify ` command to
modify the order of looking up the name service sources, for an existing
name service switch entry.

== Parameters

``-vserver <vserver name>`` - Vserver::
Vserver Use this parameter to specify the Vserver on which to modify the
name service switch entry. A data Vserver or admin Vserver can be
specified.

``-database {hosts|group|passwd|netgroup|namemap}`` - Name Service Switch
Database::
Name Service Switch Database Use this parameter to specify the name
service database, of the given Vserver,for which to modify the name
service switch entry. Following are the possible values for this
parameter:
+
  *  hosts
  *  group
  *  passwd
  *  netgroup
  *  namemap
```

```
``[-sources {files|dns|ldap|nis}]`` - Name Service Source Order::
Name Service Source Order Use this parameter to specify the name service
sources and the order in whichlook up for the specified Vserver and name
service database. Each name service source in the list for this parameter
must be one of the following:
+
  *  files
  *  dns
  *  ldap
  *  nis


+
Separate multiple sources with commas.
+
For each database specified with the `-database` parameter, one or more
sources must be specified. The valid sources for each database type are
shown in the following table:
```

```
+----------+-------------------+
| Database |  Valid Sources    |
|----------+-------------------|
| hosts    |  files, dns       |
| group    |  files, nis, ldap |
| passwd   |  files, nis, ldap |
| netgroup |  files, nis, ldap |
| namemap  |  files, ldap      |
+----------+-------------------+
```

```
+
NOTE: If "files" is not specified as the default source for the "passwd"
or "group" database, ensure that default user and group entries for the
'passwd' and 'group' respectively are present in the source configured.
Default entries for "passwd" database: nobody, pcuser, root, sshd, toor,
daemon, operator, bin, tty, kmem, games, news, man, smmsp, mailnull, bind,
proxy, uucp, pop, www, admin, diag, autosupport. Default entries for
"group" database: wheel, daemon, kmem, sys, tty, operator, mail, bin,
news, man, games, ftp, staff, sshd, smmsp, mailnull, guest, bind, proxy,
authpf, _pflogd, _dhcp, uucp, dialer, network, audit, www, antivirus,
nogroup, nobody.
+


== Examples

The following example modifies the name service source ordering for the
hosts database on a Vserver named vs0. The order of looking up the sources
is changed to only DNS.
```

cluster1::> vserver services name-service ns-switch modify -vserver vs0 -database hosts -sources dns

```
The following example modifies the name service source ordering for the
passwd database on a Vserver named vs1. The order of looking up the
sources is changed to LDAP followed by NIS.
```

cluster1::> vserver services name-service ns-switch modify -vserver vs1 -database passwd -sources ldap,nis

```
[[IDd8ed1b93862b80e6fd4b89024d383a22]]
= vserver services name-service ns-switch show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Name Service Switch configuration

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the ` vserver services name-service ns-switch show ` command to
```

```
display information about one or more name service switch entries. A name
service switch entry provides information about the order of looking up
the name service sources, for a Vserver and name service database.

== Parameters

{ ``[-fields <fieldname>,...]``::
 If you specify the -fields [fieldname], ... parameter, the command only
displays the fields that you specify.

| ``[-instance  ]`` }::
If you specify the -instance parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
Vserver Use this parameter to display only the name service switch entries
for the Vserver you specify. A data Vserver or admin Vserver can be
specified.

``[-database {hosts|group|passwd|netgroup|namemap}]`` - Name Service
Switch Database::
Name Service Switch Database Use this parameter to display only the name
service switch entries of the name service database type you specify.
Following are the possible values for this parameter:
+
  *  hosts
  *  group
  *  passwd
  *  netgroup
  *  name_map



``[-sources {files|dns|ldap|nis}]`` - Name Service Source Order::
Name Service Source Order Use this parameter to display only name service
switch entries with the specified name service source order. Each name
service source in the list for this parameter must be one of the
following:
+
  *  files
  *  dns
  *  ldap
  *  nis


+
Separate multiple sources with commas.
```

```
== Examples

The following example shows the output of the ` vserver services name-
service ns-switch show ` command.
```

cluster1::> ` vserver services name-service ns-switch show `

```
                                    Source
Vserver         Database              Order
--------------- ------------          ---------
vs0             hosts                 files,
                                      dns
vs1             passwd                files,
                                      ldap, nis
2 entries were displayed.
```

```
[[ID838880f11f4d821cce8114571593bc47]]
= vserver services name-service unix-group adduser
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add a user to a local UNIX group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-group adduser` command adds a user
to a local UNIX group.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver location of the local UNIX group
to which the user is added.

``-name <text>`` - Group Name::
Use this parameter to specify the local UNIX group to which to add the
user.
```

```
``-username <text>`` - Name of User::
Use this parameter to specify the user name to add to the local UNIX
group.

``[-skip-name-validation {true|false}]`` - Skip Name Validation::
By default, ONTAP validates the name to ensure it complies with the
following rules:
+
  * The name contains only these valid characters: 0 through 9, A through
Z, a through z, "_", ".", and "-".
  * The name does not start with the character "-".
  * The name does not contain "$" except as the last character.


+
If the parameter is set to ``_true_`` , the name validation is skipped.

== Examples

The following example adds a user named tsmith to a local UNIX group named
sales on a Vserver named vs0:
```

cluster1::> vserver services name-service unix-group adduser -vserver vs0 -name sales -username tsmith

```
[[ID69012d7aaa40f059c9f7fd720dfeff66]]
= vserver services name-service unix-group create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a local UNIX group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-group create` command creates a
local UNIX group on a Vserver. Use a local UNIX group for Windows-to-UNIX
and UNIX-to-Windows group mappings.

== Parameters
```

```
``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver on which to create the local
UNIX group.

``-name <text>`` - Group Name::
Use this parameter to specify the name of the group to create.

``-id <integer>`` - Group ID::
Use this parameter to specify an ID number for the group.

``[-skip-name-validation {true|false}]`` - Skip Name Validation::
By default, ONTAP validates the name to ensure it complies with the
following rules:
+
  * The name contains only valid characters: 0 through 9, A through Z, a
through z, "_", ".", and "-"
  * The name does not start with "-"
  * The name does not contain "$" except as the last character


+
If the parameter is set to ``_true_`` , the name validation is skipped.

== Examples

The following example creates a group named sales on a Vserver named vs0.
The group has the ID 94.
```

cluster1::> vserver services name-service unix-group create -vserver vs0 -name sales -id 94

```
[[ID8447ae234ed314bd5497e1cc67b2e496]]
= vserver services name-service unix-group delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a local UNIX group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-group delete` command deletes a
local UNIX group from a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver location of the local UNIX group
to delete.

``-name <text>`` - Group Name::
Use this parameter to specify the local UNIX group to delete.

== Examples

The following example deletes a local UNIX group named testgroup from a
Vserver named vs0:
```

cluster1::> vserver services name-service unix-group delete -vserver vs0 -name testgroup

```
[[IDf093f0bb4307da65c5f85acf552faad5]]
= vserver services name-service unix-group deluser
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a user from a local UNIX group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-group deluser` command removes a
user from a local UNIX group.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver location of the local UNIX group
from which the user is removed.

``-name <text>`` - Group Name::
Use this parameter to specify the local UNIX group from which to remove
the user.

``-username <text>`` - Name of User::
Use this parameter to specify the user name to remove from the local UNIX
group.

== Examples

The following example removes a user named testuser from a local UNIX
group named sales on a Vserver named vs0:
```

cluster1::> vserver services name-service unix-group deluser -vserver vs0 -name eng -username testuser

```
[[ID4472cd23567fd614ca9d31b226b0968d]]
= vserver services name-service unix-group load-from-uri
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Load one or more local UNIX groups from a URI

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-group load-from-uri` command loads
UNIX groups from a universal resource identifier (URI). The URI must
contain group information in the UNIX /etc/group format:

``_group_name_`` :``_password_`` :``_group_ID_``
:``_comma_separated_list_of_users_``

The command discards the value of the ``_password_`` field.

== Parameters

``-vserver <vserver>`` - Vserver::
Use this parameter to specify the Vserver on which to locate the local
UNIX groups.

``-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}`` - URI to
Load From::
Use this parameter to specify the URI from which the command loads group
information. The URI scheme must be either ftp(s) or http(s)

``[-overwrite {true|false}]`` - Overwrite Entries::
Use this parameter with the value `true` to specify that group information
loaded from the URI should overwrite existing group information. The
default value is `false` , specifying that group information loaded from
the URI should not overwrite existing group information.

``[-skip-name-validation {true|false}]`` - Skip Name Validation::
By default, ONTAP validates the name to ensure it complies with the
following rules:
+
  * The name contains only valid characters: 0 through 9, A through Z, a
through z, "_", ".", and "-"
  * The name does not start with "-"
  * The name does not contain "$" except as the last character



+
If the parameter is set to ``_true_`` , the name validation is skipped.
```

```
``[-foreground {true|false}]`` - Load Unix Groups file in the Foreground::
If this parameter is set to ``_false_`` , the operation runs as a job in
the background. Otherwise, the command does not return until the operation
is complete. The default value is ``_true_`` .

== Examples

The following example loads group information from the URI
ftp://ftp.example.com/groups onto a Vserver named vs0:
```

cluster1::> vserver services name-service unix-group load-from-uri -vserver vs0 -uri ftp://ftp.example.com/groups

```
[[ID23b1839d9fa5707d25d5ca9c74454aef]]
= vserver services name-service unix-group modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a local UNIX group

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

Use the `vserver services name-service unix-group modify` command to
modify a local UNIX group's group ID.

== Parameters

``-vserver <vserver name>`` - Vserver::
Use this parameter to specify the Vserver location of the local UNIX group
to modify.

``-name <text>`` - Group Name::
Use this parameter to specify the name of the group to modify.

``[-id <integer>]`` - Group ID::
Use this parameter to specify an ID number for the group.

== Examples

The following example changes a local UNIX group named sales on a Vserver
named vs0 to have the group ID 100:
```

cluster1::> vserver services name-service unix-group modify -vserver vs0 -group sales -id 100

```
[[IDd3f72a8e2cc40d48f3200cff199a0086]]
= vserver services name-service unix-group show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Display local UNIX groups

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-group show` command displays
information about local UNIX groups.

== 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.

| ``[-members  ]``::
Use this parameter to display the list of users in each local UNIX group.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver::
Use this parameter with the `-name` parameter to display information only
about the local UNIX group you specify. Use this parameter without `-name`
to display information only about the local UNIX groups that are located
on the specified Vserver.

``[-name <text>]`` - Group Name::
Use this parameter with the `-vserver` parameter to display information
only about the local UNIX group you specify. Use this parameter without `-
vserver` to display information only about the local UNIX groups that
match the name you specify.

``[-id <integer>]`` - Group ID::
Use this parameter to display information only about the local UNIX group
that has the ID you specify.

``[-users <text>,...]`` - Users::
Use this parameter to display information only about the local UNIX groups
that include the user names you specify.

== Examples

The following example displays information about all local UNIX groups,
```

```
including lists of their users:
```

cluster1::> vserver services name-service unix-group show -members
Vserver Name ID

vs0 dev 44
Users: admin, jdoe, tsmith

vs0 sales 12
Users: admin, guest, pjones

vs1 testgroup 13
Users: admin, root, testuser

vs1 users 100
Users: admin, jdoe, pjones, tsmith

```
[[IDb7854b94f384f46145fe39bc186b8394]]
= vserver services name-service unix-group file show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local UNIX groups file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service unix-group file show` command displays
information about local UNIX groups. It displays the content as it is from
the actual UNIX group file which resides in the mroot volume.

== Parameters

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
local UNIX group or groups that are located on the specified Vserver.

``[-search-string <text>]`` - Pattern to be searched [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter and the `-vserver` parameter, the command
only displays information from the UNIX group file which matches the
specified parameter.

== Examples

The following example displays information about all local UNIX groups
belonging to a specific Vserver:
```

cluster1::> vserver services name-service unix-group file show -vserver vs0
Line No File content
--- ---
1 daemon:*:1:
2 nobody:*:65535:
3 pcuser:*:65534:
4 root:*:0:

```
[[ID03764e58cc15de90eb35e9e37e148ec7]]
= vserver services name-service unix-group file status
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local Unix Groups file status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service unix-group file status` command
displays the status of local UNIX group file across a cluster. This
enables you to verify that UNIX group files are consistent across all
nodes that back a Vserver into which UNIX group files have been loaded.

The command displays the following information:

* Vserver name
* Node name
* Load time for the UNIX group file
* Hash value of the UNIX group file
* Hash value of the UNIX group database file
* Hash value of the UNIX group byuser database file
* File size of the UNIX group file



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the UNIX group status
only for the specified Vserver.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
```

```
advanced)#::
If you specify this parameter, the command displays the UNIX group status
only for the specified node.

``[-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Load Time [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the UNIX group file that were loaded at the specified time. Specify time
in the format MM/DD/YYYY HH:MM:SS. Note that the load time stamps for
identical files are different on different nodes, because each node
downloads the definitions from the source URI individually.

``[-hashvalue <text>]`` - Hash Value [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, command displays the status only for the
UNIX group files that have the specified hash value. Note that the primary
purpose of the command is to verify that the files on all nodes have the
same hash value, so querying on a specific hash value is not useful in
most cases.

``[-hashvalue-db-grp <text>]`` - Hash Value DB [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, command displays the status only for the
UNIX group files that have the specified hash value for the UNIX group
database. Note that the primary purpose of the command is to verify that
the files on all nodes have the same hash value, so querying on a specific
hash value is not useful in most cases.

``[-hashvalue-db <text>]`` - Hash Value byuser DB [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the UNIX group files that have the specified hash value for the UNIX group
byuser database. Note that the primary purpose of the command is to verify
that the files on all nodes have the same hash value for UNIX group
database.

``[-filesize {<integer>[KB|MB|GB|TB|PB]}]`` - File Size [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the UNIX group files that have the specified file size. Note that the
primary purpose of the command is to verify that the files on all nodes
have the same file size, so querying on a specific file size is not useful
in most cases.

== Examples
```

```
The following example displays the UNIX group file status for all
Vservers:
```

cluster1::*> vserver services name-service unix-group file status -instance
Vserver: vs1
Node: node1
Load Time: 8/9/2016 19:56:25
Hash Value: 835c7f530fb76f96c3bca00e380d36b7
Hash Value DB: e6cb38ec1396a280c0d2b77e3a84eda2
Hash Value byuser DB: 913a182a72aa1872495be398ebb2cd23
File Size: 58B
Vserver: vs2
Node: node1
Load Time: 8/9/2016 20:15:40
Hash Value: c0d2b77e3a84eda2e6cb38ec1396a280
Hash Value DB: 009321eddb45611e95d9f7f277ec0621
Hash Value byuser DB: 659321eddb45611e95d9f7f277ec0621
File Size: 2.3MB
2 entries were displayed.

```
[[ID53266baba6939dbd3759033a365f53d5]]
= vserver services name-service unix-group max-limit modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change Configuration Limits for UNIX-Group

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service unix-group max-limit modify` command
enables you to modify maximum UNIX groups and group-members that can be
configured on the system. This allows you to set certain limits to prevent
performance issues due to service configurations using excessive
resources.

== Parameters

``[-limit <integer>]`` - System Limit [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the maximum limit that you want to set for unix-
group. The default setting for the limit is 32768. The supported range of
values for this parameter is 0 to 65536.

== Examples

The following example modifies the system-wide limit of the total number
of UNIX groups and members that can be configured on the cluster.
```

vserver services name-service unix-group max-limit modify -limit 33792

```
[[ID154e513be1fcfdc0ba5c5611d4f007b9]]
= vserver services name-service unix-group max-limit show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Configuration Limits for UNIX-Group

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service unix-group max-limit show` command
displays information on UNIX group limits that are configurable with
xref:{relative_path}vserver-services-name-service-unix-group-max-limit-
modify.html[vserver services name-service unix-group max-limit modify]
command. The output will show the following:

* Limit: The configured limit on the total number of UNIX groups and group
members configurable.
* Current Count: Total number of current entries for UNIX groups and group
members.



== Examples

The following example shows the limits and total number of current entries
for UNIX group configuration:
```

cluster1::> vserver services name-service unix-group max-limit show
(vserver services name-service unix-group max-limit show)
Limit Current Count
---------- ----------
400 3

```
== Related Links

* xref:{relative_path}vserver-services-name-service-unix-group-max-limit-
modify.html[vserver services name-service unix-group max-limit modify]
```

```
[[ID79562ee5527433515aa9c5dd14fa877a]]
= vserver services name-service unix-user create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a local UNIX user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-user create` command creates a
local UNIX user on a Vserver. You can use local UNIX users for Windows-to-
UNIX and UNIX-to-Windows name mappings.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which you want to create the local
unix user.

``-user <text>`` - User Name::
This parameter specifies the user account that you want to create.

``-id <integer>`` - User ID::
This parameter specifies an ID number for the user.

``-primary-gid <integer>`` - Primary Group ID::
This parameter specifies the ID number of the user's primary group.

``[-full-name <text>]`` - User's Full Name::
This parameter specifies the user's full name.

``[-skip-name-validation {true|false}]`` - Skip Name Validation::
By default, ONTAP validates the name to ensure it complies with the
following rules:
+
  * The name contains only valid characters: 0 through 9, A through Z, a
through z, "_", ".", and "-"
  * The name does not start with "-"
  * The name does not contain "$" except as the last character
```

```
+
If the parameter is set to ``_true_`` , the name validation is skipped.

== Examples

The following example creates a local UNIX user named tsmith on a Vserver
named vs0. The user has the ID 4219 and the primary group ID 100. The
user's full name is Tom Smith.
```

vs1::> vserver services name-service unix-user create -vserver vs0 -user tsmith -id 4219 -primary-gid 100 -full -name "Tom Smith"

```
[[IDfab9a2726b0e17df52f3984f0d545b32]]
= vserver services name-service unix-user delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a local UNIX user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-user delete` command deletes a
local UNIX user from a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the local UNIX user is
located.

``-user <text>`` - User Name::
This parameter specifies the user that you want to delete.

== Examples

The following example deletes a local UNIX user named testuser from a
Vserver named vs0:
```

vs1::> vserver services name-service unix-user delete -vserver vs0 -user testuser

```
[[ID4c0c97894b0a0a99490148efc13446cd]]
= vserver services name-service unix-user load-from-uri
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Load one or more local UNIX users from a URI

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-user load-from-uri` command loads
one or more UNIX users from a universal resource identifier (URI). The URI
must contain user information in the UNIX /etc/passwd format:
``_user_name_`` :``_password_`` :``_user_ID_`` :``_group_ID_``
:``_full_name_`` :``_home_directory_`` :``_shell_`` . The command discards
the value of the ``_password_`` field and of the fields after the
``_full_name_`` field ( ``_home_directory_`` and ``_shell_`` ).

== Parameters

``-vserver <vserver>`` - Vserver::
This specifies the Vserver on which the local UNIX user or users are to be
located.

``-uri {scheme://(hostname|IPv4 Address|'['IPv6 Address']')...}`` - URI to
Load From::
This specifies the URI from which user information is to be loaded. The
URI scheme must be either ftp(s) or http(s).

``[-overwrite {true|false}]`` - Overwrite Entries::
This optionally specifies whether user information from the URI overwrites
existing user information. The default setting is `false` .

``[-skip-name-validation {true|false}]`` - Skip Name Validation::
By default, ONTAP validates the name to ensure it complies with the
following rules:
+
  * The name contains only valid characters: 0 through 9, A through Z, a
through z, "_", ".", and "-"
```

```
  * The name does not start with "-"
  * The name does not contain "$" except as the last character


+
If the parameter is set to ``_true_`` , the name validation is skipped.

``[-foreground {true|false}]`` - Load Unix Users file in the Foreground::
If this parameter is set to ``_false_`` , the operation runs as a job in
the background. Otherwise, the command does not return until the operation
is complete. The default value is ``_true_`` .

== Examples

The following example loads user information from the URI
ftp://ftp.example.com/users onto a Vserver named vs0:
```

node::> vserver services name-service unix-user load-from-uri -vserver vs0 -uri ftp://ftp.example.com/users

```
[[ID3f44192af3a3d5ee3124295d215c4f99]]
= vserver services name-service unix-user modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a local UNIX user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-user modify` command modifies a
local UNIX user's ID, primary group ID, or full name.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the Vserver on which the local UNIX user is
located.

``-user <text>`` - User Name::
This parameter specifies the user account that you want to modify.

``[-id <integer>]`` - User ID::
This optional parameter specifies an ID number for the user.

``[-primary-gid <integer>]`` - Primary Group ID::
This optional parameter specifies the ID number of the user's primary
group.

``[-full-name <text>]`` - User's Full Name::
This optional parameter specifies the user's full name.

== Examples

The following example modifies the local UNIX user named pjones on a
Vserver named vs0. The user's primary group ID is changed to 100 and the
user's full name is Peter Jones.
```

vs1::> vserver services name-service unix-user modify -vserver vs0 -user pjones -primary-gid 100 -full-name "Peter Jones"

```
[[ID4ee0723c335fe16b351e07642d844349]]
= vserver services name-service unix-user show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local UNIX users

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services name-service unix-user show` command displays
information about local UNIX users. The command output depends on the
parameter or parameters specified with the command. If you do not specify
any parameters, the command displays the following information about all
local UNIX users:

* Vserver name
* User name
* User ID
* Primary group ID
* Full name



== Parameters

{ ``[-fields <fieldname>,...]``::
If you specify the `-fields` parameter, the command only displays the
fields that you specify.

| ``[-instance  ]`` }::
If you specify the `-instance` parameter, the command displays detailed
information about all entries.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter and the `-user` parameter, the command
displays information only about the specified local UNIX user. If you
specify this parameter by itself, the command displays information only
about the local UNIX user or users that are located on the specified
Vserver.

``[-user <text>]`` - User Name::
```

```
If you specify this parameter and the `-vserver` parameter, the command
displays information only about the specified local UNIX user. If you
specify this parameter by itself, the command displays information only
about the local UNIX user or users that have the specified name.

``[-id <integer>]`` - User ID::
If you specify this parameter, the command displays information only about
the local UNIX user that has the specified ID.

``[-primary-gid <integer>]`` - Primary Group ID::
If you specify this parameter, the command displays information only about
the local UNIX user or users that have the specified primary group ID.

``[-full-name <text>]`` - User's Full Name::
If you specify this parameter, the command displays information only about
the local UNIX user or users that match the specified name.

== Examples

The following example displays information about all local UNIX users:
```

vs1::> vserver services name-service unix-user show
User User Group Full
Vserver Name ID ID Name
---------- ---------- ------ ------ ----------
vs0 admin 100 100 administrator
vs0 guest 1000 100 guest
vs0 jdoe 4673 100 Jane Doe
vs0 monitor 2000 100 monitor
vs0 pjones 4236 100 Peter Jones
vs0 root 10 100 root
vs0 tsmith 3289 100 Tom Smith

```
[[ID3d7be186702fa48408188151891f8c8b]]
= vserver services name-service unix-user file show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local UNIX users file

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service unix-user file show` command displays
information about local UNIX users. It displays the content as it is from
the actual UNIX user file which resides in the mroot volume.

== Parameters

``-vserver <vserver>`` - Vserver [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information about the
local UNIX user or users that are located on the specified Vserver.

``[-search-string <text>]`` - Pattern to be searched [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter and the `-vserver` parameter, the command
only displays information from the UNIX user file which matches the
specified parameter.

== Examples

The following example displays information about all local UNIX users
belonging to a specific Vserver:
```

cluster1::> vserver services name-service unix-user file show -vserver vs0
Line No File content
---------- ----------
1 nobody:*:65535:65535::::::
2 pcuser:*:65534:65534::::::
3 root:*:0:1::::::

```
[[IDadf41ef26578516b474cb8a8085bd926]]
= vserver services name-service unix-user file status
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display local Unix Users file status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver services name-service unix-user file status` command displays
the status of local UNIX user file across a cluster. This enables you to
verify that UNIX user files are consistent across all nodes that back a
Vserver into which UNIX user files have been loaded.

The command displays the following information:

* Vserver name
* Node name
* Load time for the UNIX user file
* Hash value of the UNIX user file
* Hash value of the UNIX user database file
* File size of the UNIX user file



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the UNIX user status
only for the specified Vserver.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
```

```
If you specify this parameter, the command displays the UNIX user status
only for the specified node.

``[-timestamp <MM/DD/YYYY HH:MM:SS>]`` - Load Time [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the UNIX user file that were loaded at the specified time. Specify time in
the format MM/DD/YYYY HH:MM:SS. Note that the load time stamps for
identical files are different on different nodes, because each node
downloads the definitions from the source URI individually.

``[-hashvalue <text>]`` - Hash Value [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, command displays the status only for the
UNIX user files that have the specified hash value. Note that the primary
purpose of the command is to verify that the files on all nodes have the
same hash value, so querying on a specific hash value is not useful in
most cases.

``[-hashvalue-db <text>]`` - Hash Value DB [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays the status only for
the UNIX user files that have the specified hash value for the UNIX user
database. Note that the primary purpose of the command is to verify that
the files on all nodes have the same hash value for UNIX user database.

``[-filesize {<integer>[KB|MB|GB|TB|PB]}]`` - File Size [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays the status only for
the UNIX user files that have the specified file size. Note that the
primary purpose of the command is to verify that the files on all nodes
have the same file size, so querying on a specific file size is not useful
in most cases.

== Examples

The following example displays the UNIX user file status for all Vservers:
```

cluster1::*> vserver services name-service unix-user file status
Vserver Node Load Time Hash Value Hash Value DB File Size
--------- ------- ------------------- ---------------------------------- ---------------------------------- ----------
vs1
node1 5/20/2016 16:04:55 e6cb38ec1396a280c0d2b77e3a84eda2 913a182a72aa1872495be398ebb2cd23
1.00KB
node2 5/20/2016 16:04:53 e6cb38ec1396a280c0d2b77e3a84eda2 913a182a72aa1872495be398ebb2cd23
1.00KB
vs2

node1 5/20/2016 16:06:26 c0d2b77e3a84eda2e6cb38ec1396a280 009321eddb45611e95d9f7f277ec0621 2.3MB
node2 5/20/2016 16:06:27 c0d2b77e3a84eda2e6cb38ec1396a280 009321eddb45611e95d9f7f277ec0621 2.3MB
4 entries were displayed.

```
[[IDb50f5ba2eee4e746cc3845daa7e18237]]
= vserver services name-service unix-user max-limit modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Change Configuration Limits for UNIX-User

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service unix-user max-limit modify` command
enables you to modify maximum UNIX users that can be configured on the
system. This allows you to set certain limits to prevent performance
issues due to service configurations using excessive resources.

== Parameters

``[-limit <integer>]`` - System Limit [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the maximum limit that you want to set for unix-
user. The default setting for the limit is 32768. The supported range of
values for this parameter is 0 to 65536.

== Examples

The following example modifies the system-wide limit of the total number
of UNIX users that can be configured on the cluster.
```

vserver services name-service unix-user max-limit modify -limit 33792

```
[[ID7d50fa8a352668f7aed1ba59b2c9feef]]
= vserver services name-service unix-user max-limit show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Configuration Limits for UNIX-User

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service unix-user max-limit show` command
displays information on UNIX user limits that are configurable with
xref:{relative_path}vserver-services-name-service-unix-user-max-limit-
modify.html[vserver services name-service unix-user max-limit modify]
command. The output will show the following:

* Limit: The configured limit on the total number of UNIX users
configurable.
* Current Count: Total number of current entries for UNIX users
configuration.



== Examples

The following example shows the limits and total number of current entries
for UNIX user configuration:
```

cluster1::> vserver services name-service unix-user max-limit show
(vserver services name-service unix-user max-limit show)
Limit Current Count
--- ---
400 3

```
== Related Links

* xref:{relative_path}vserver-services-name-service-unix-user-max-limit-
modify.html[vserver services name-service unix-user max-limit modify]




[[ID5ae1e82460c33206d4eaf7f3e7efc67c]]
= vserver services name-service ypbind start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start ypbind

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service ypbind start` starts the ypbind. NIS
creation will fail if ypbind is stopped. This command starts ypbind on all
the nodes in a cluster and is presistent across node reboots.

== Examples

The following example starts ypbind:
```

vs1::> vserver services name-service ypbind start

```
[[ID9ff0e6235efbbb36ff0764d6a1247c47]]
= vserver services name-service ypbind status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Current ypbind status

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service ypbind status` displays whether the
ypbind is running or stopped. The ypbind service starts by default when a
node comes up unless the service is in the stopped state, irrespective of
whether NIS is configured.

== Examples

The following example displays ypbind status:
```

vs1::> vserver services name-service ypbind status
Status: Running

```
[[ID55b9e5456336dcec8b1b3c237a8ab1e3]]
= vserver services name-service ypbind stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop ypbind

*Availability:* This command is available to _cluster_ administrators at
the _advanced_ privilege level.

== Description

The `vserver services name-service ypbind stop` stops the ypbind. Command
fails if NIS entries are present. This command stops ypbind on all the
nodes in a cluster and is persistent across node reboots.

== Examples

The following example stops ypbind:
```

vs1::> vserver services name-service ypbind stop

```
[[ID10248259c1b0467b47ab97fc34b3ca92]]
= vserver services ndmp generate-password
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Generates NDMP password for a user

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command is used to generate NDMP password for a given user in the
specified Vserver context. The generated NDMP password is based on the
user's login password. For this reason regenerate it whenever the user's
login password changes. This command fails if a user does not exist for
the Vserver.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Specify the Vserver context for which password is to be generated.

``[-user <text>]`` - User::
Specify the user name for which the NDMP password needs to be generated.

``[-password <text>]`` - Password::
The generated NDMP password string that is used for authentication.

== Examples

The following example shows the usage this command to generate NDMP
password for a user belonging to a specific Vserver:
```

cluster1::> vserver services ndmp generate-password -vserver vserver1 -user user1
Vserver: vserver1
User: user1
Password: a9cCCUp32yjGmBiD

```
[[ID4e458db40fd4e60e403f68d4b429c9b0]]
= vserver services ndmp kill-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Kill all NDMP sessions

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command terminates all NDMP sessions on a particular Vserver in the
cluster.

== Parameters

``-vserver <vserver name>`` - Vserver::
Specifies the Vserver name in which all NDMP sessions that are to be
terminated are running.

== Examples

The following example shows how all NDMP sessions on the Vserver named
vserver1 can be terminated:
```

cluster1::> vserver services ndmp kill-all -vserver vserver1

```
[[ID90efc168a23de15b28dcace0795d9095]]
= vserver services ndmp kill
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Kill the specified NDMP session

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command terminates a specific NDMP session on a particular Vserver in
the cluster.

== Parameters

``<text>`` - Session Identifier::
Session ID of the NDMP session. A session-id is a string used to identify
a particular NDMP session.

== Examples

The following example shows how a specific NDMP session on the Vserver
named vserver1 can be terminated:
```

cluster1::> vserver services ndmp kill 1000:8002 -vserver vserver1

```
[[IDbe04bc703ff033d9d0782a1e7e7ee25d]]
= vserver services ndmp modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify  NDMP Properties

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description
```

```
This command is used to change NDMP options on Vservers.

One or more of the options specified in the parameters section can be
modified for a specific Vserver, by this command. A short description of
each of the options is provided in the parameters section.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the name of the Vserver.

``[-ignore-ctime-enabled {true|false}]`` - Ignore Ctime::
This option, when ``_true_`` , allows user to exclude files with ctime
changed from storage system' incremental dumps since other processes like
virus scanning often alter the ctime of files. When this option is
``_false_`` , backup on the Vserver will include all files with a change
or modified time later then the last dump in the previous level dump. The
default value is ``_false_`` . This option is persistent across reboots.
+
Most WIN32 APIs are often unaware of the "last changed time", ctime, they
often incorrectly set a later time for files, causing these files to be
included in the Vserver's incremental dumps, making the incremental dump
very large. This is partially defying the purpose of having incremental
dumps, since one uses incremental dumps to speed up the backup by only
dumping files that were truly changed since the last backup.
+
The `-option-value` for this parameter should be true/false.

``[-offset-map-enable {true|false}]`` - Enable Offset Map::
This option is used to enable or disable generation of the inode offset
map during NDMP based dump backups. The offset map is required to perform
Enhanced Direct Access Restore (DAR) on the backup data. Enhanced DAR
provides support for directory DAR and DAR of files with NT streams. The
default value for this option is ``_true_`` . This option is persistent
across reboots.
+
The `-option-value` for this parameter should be true/false.

``[-tcpnodelay {true|false}]`` - Enable TCP Nodelay::
Enables/Disables the TCPNODELAY configuration parameter for the socket
between the Vserver and the DMA. When set to true, the Nagle algorithm is
disabled and small packets are sent immediately rather than held and
bundled with other small packets. This optimizes the system for response
time rather than throughput.
+
```

```
This option becomes active when the next NDMP session starts. Existing
sessions are unaffected. The default value for this option is ``_false_``
. This option is persistent across reboots.
+
The `-option-value` for this parameter should be true/false.

``[-tcpwinsize <integer>]`` - TCP Window Size::
This option can be used to change the TCP buffer size of the NDMP data
connection.The minimum and maximum values are 8192(8K) and
7,631,441(7.2M), respectively. The default value for this option is
32768(32K).
+
This option is persistent across reboots.
+
The `-option-value` for this parameter should be a number between 8192(8K)
and 7,631,441(7.2M).

``[-data-port-range <text>]`` - Data Port Range::
This option allows administrators to specify a port range on which the
NDMP server can listen for data connections.
+
The format of this option is ``_start_port - end_port_`` start_port,
end_port can have values between [1024-65535]; start_port must be lesser
than or equal to end_port. If a valid range is specified, NDMP uses a port
within that range to listen for data connections. A listen request fails
if no ports in the specified range are free.
+
This option is modifiable only from the admin Vserver context and the said
option is applicable for all the data Vservers and the admin Vserver. For
example, if the value of the above option is set with 2000-3000, the same
value will be applicable throughout the cluster. The value ``_all_``
implies that any available port can be used to listen for data
connections. The default value for this option is ``_all_`` . This option
is persistent across reboots.
+
The `-option-value` for this option should be in the format {<start_port>-
<end port> | all }- where start_port, end_port can have values between
[1024-65535]; start_port must be lesser than or equal to end_port.

``[-backup-log-enable {true|false}]`` - Enable Backup Log::
Backup logging captures important events during dump/restore and records
them in /mroot/etc/log/backup on the root volume. The option allows users
to enable or disable this feature. The default value for this option is
``_true_`` . This option is persistent across reboots.
+
The `-option-value` for this parameter should be true/false.
```

```
``[-per-qtree-exclude-enable {true|false}]`` - Enable per Qtree
Exclusion::
If this option is ``_true_`` , users can specify exclude list on a per
qtree basis to be excluded from backup. This exclude list will override
any values already present due to 'EXCLUDE' environment variable . The
user can specify the exclusion list through a .exclude_list file which
resides at the root of the qtree. The exclusion list can be a list of
files or files that match a specified pattern. The default value for this
option is ``_false_`` . This option is persistent across reboots.
+
The `-option-value` for this parameter should be true/false.

``[-authtype <NDMP Authentication types>,...]`` - Authentication Type::
Allows the administrator to choose the authentication method. NDMP
supports three authentication types: challenge, plaintext and
plaintext_sso. The plaintext_sso authentication type is mutually exclusive
with the other authentication types. By setting the authentication type as
plaintext_sso, the actual password for the user can be used to
authenticate instead of having to generate an NDMP specific password. The
default of this option is ``_challenge_`` . This option is persistent
across reboots.
+
The `-option-value` for this parameter can be {challenge | plaintext |
plaintext_sso | challenge, plaintext | plaintext, challenge}.

``[-debug-enable {true|false}]`` - Enable Debug [.advaced-
priv]#(privilege: advanced)#::
This option enables debug logging for NDMP. Debug messages will be logged
to the ndmpd log file /mroot/etc/log/mlog/ndmpd.log . The default value
for this option is ``_false_`` .This option is persistent across reboots.
+
The `-option-value` for this parameter should be true/false.

``[-debug-filter <text>]`` - Debug Filter [.advaced-priv]#(privilege:
advanced)#::
This option controls the NDMP modules for which debug logging is to be
enabled. option-value can take five values for this option : all, none,
normal, backend or "filter-expression".
+
``_all_`` enables debug logging for all modules.
+
``_none_`` disables debug logging for all modules. It is equivalent to
``_modify -vserver vserver_name -debug-enable false_`` .
+
``_normal_`` is a shortcut option that enables debug logging for all
```

```
modules except verbose and io_loop. The equivalent filter string is all-
verbose-io_loop.
+
``_backend_`` is a short cut option that enables debug logging for all
modules except verbose, io_loop, ndmps and ndmpd. The equivalent filter
string is all-verbose-io_loop-ndmps-ndmpp.
+
``_(filter-expression)_`` is a combination of one or more modules for
which debug logs needs to be enabled. Multiple module names can be
combined using following operators :
+
  * - to remove the given module from the list of specified modules in the
filter string. For example the filter all-ndmpp will enable debug logging
for all modules but not ndmpp.
  * ^ to add the given module or modules to the list of modules specified
in the filter string. For example the filter ndmpp^mover^data will enable
debug logging for ndmpp, mover and data.

+
The possible module names and a brief description is given below:-
```

```
+--------------+-------------------------------------+
| Modules      | Description                         |
+--------------+-------------------------------------+
| verbose      | verbose message                     |
| io           | I/O process loop                    |
| io_loop      | I/O process loop verbose messages   |
| ndmps        | NDMP service                        |
| ndmpp        | NDMP Protocol                       |
| rpc          | General RPC service                 |
| fdc_rpc      | RPC to FC driver service            |
| auth         | Authentication                      |
| mover        | NDMP MOVER (tape I/O)               |
| data         | NDMP DATA (backup/restore)          |
| scsi         | NDMP SCSI (robot/tape ops)          |
| bkup_rpc     | RPC to Backup service client        |
| bkup_rpc_s   | RPC to Backup service server        |
| conf         | Debug configure/reconfigure         |
| dblade       | Dblade specific messages            |
| timer        | NDMP server timeout messages        |
| vldb         | VLDB service                        |
| smf          | SMF Gateway messages                |
| common       | NDMP common state                   |
| ext          | NDMP extensions messages            |
| ndmprpc      | NDMP Mhost RPC server               |
+--------------+-------------------------------------+
```

```
+
The default value for this option is ``_none_`` . This option is
persistent across reboots.
+
The `-option-value` for this parameter can be {all | none | normal |
backend |'filter-expression'}.

``[-dump-logical-find <text>]`` - Enable Logical Find for Dump [.advaced-
priv]#(privilege: advanced)#::
This option specifies whether to follow inode-file walk or tree walk for
phase I of the dump. Choosing inode-file walk or tree walk affects the
performance of the dump. This option can take following values:
+
If ``_default_`` is specified, then level 0 and incremental volume as well
as qtree dumps will use inode walk. All the subtree dumps will use tree
walk.
+
If ``_always_`` is specified, all dumps will follow treewalk.
```

```
+
A ``_comma-separated_`` list of values in any combination from the
following list:
+
  * vol_baseline: Level 0 full volume backup will follow treewalk.
  * vol_incr: Incremental full volume backup will follow treewalk.
  * qtree_baseline: Level 0 qtree backup will follow treewalk.
  * qtree_incr: Incremental qtree backup will follow treewalk.



+
The default value for this option is ``_default_`` . This option is
persistent across reboots.
+
The `-option-value` for this parameter could be {default | always |
'vol_baseline' | 'vol_baseline,qtree_baseline' | ...}.

``[-abort-on-disk-error {true|false}]`` - Enable Abort on Disk Error
[.advaced-priv]#(privilege: advanced)#::
If this option is ``_true_`` , dump will abort the backup operation on
detection of irrecoverable data blocks in user files. If this option is
``_false_`` , dump will proceed with backup operation - even if
irrecoverable data blocks in user files are detected. On detection of
irrecoverable data blocks, dump will send a log message to DMA and also
log an entry in /mroot/etc/log/backup file. The default value for this
option is ``_false_`` . This option is persistent across reboots.
+
The `-option-value` for this parameter should be true/false.

``[-fh-dir-retry-interval <integer>]`` - FH Throttle Value for Dir
[.advaced-priv]#(privilege: advanced)#::
NDMP protocol sends back file history information for all directories in
phase 3 of dump to DMA. In the presence of slow DMA or high latency
networks, the amount of file history being generated exceeds the amount
being consumed by the DMA. To handle a slow reader, a flow control
mechanism is now introduced where file history generation is throttled
when a DMA is slow in consuming them. The value for this option indicates
how frequently should the file history be resent if it was throttled. The
default value is 250 milliseconds. This option is persistent across
reboots.
+
The `-option-value` for this parameter should be a number.

``[-fh-node-retry-interval <integer>]`` - FH Throttle Value for Node
[.advaced-priv]#(privilege: advanced)#::
NDMP protocol sends back file history information for all files in phase 4
```

```
of dump to DMA. In the presence of slow DMA or high latency networks, the
amount of file history being generated exceeds the amount being consumed
by the DMA. To handle slow reader conditions, a flow control mechanism is
now introduced where file history generation is throttled when a DMA is
slow in consuming them. The value for this option indicates how frequently
should the file history be resent if it was throttled. The default value
is 250 milliseconds. This option is persistent across reboots.
+
The `-option-value` for this parameter should be a number.

``[-restore-vm-cache-size <integer>]`` - Restore VM File Cache Size
[.advaced-priv]#(privilege: advanced)#::
This option mandates the number of WAFL buffers pinned in memory by
various meta-files used by logical restore. The minimum and maximum values
are 4 and 1024, respectively. The default value for this option is 64.
This option is persistent across reboots.
+
Depending on the value of this option, various meta-files are assigned a
number of WAFL buffers that need to be pinned in memory.
```

```
+-----------------+---------------------------------+
| Meta-filename   |     Number of WAFL buffers to   |
|                 |      be pinned in memory        |
+-----------------+---------------------------------+
| dumpmap         | ndmpd.restore.vm_cache_size     |
| filemap         | ndmpd.restore.vm_cache_size     |
| aclfile_map     | ndmpd.restore.vm_cache_size     |
| inomap          | ndmpd.restore.vm_cache_size / 2 |
| basemap         | ndmpd.restore.vm_cache_size / 2 |
| flipmap         | ndmpd.restore.vm_cache_size / 2 |
| revmap          | ndmpd.restore.vm_cache_size / 2 |
| clrimap         | ndmpd.restore.vm_cache_size / 4 |
| mfp_for_inotab  | ndmpd.restore.vm_cache_size / 4 |
| map             | ndmpd.restore.vm_cache_size / 4 |
| offsetfile_map  | ndmpd.restore.vm_cache_size / 4 |
+-----------------+---------------------------------+
```

```
+
The `-option-value` for this parameter should be a number between 4 and
1024.

``[-enable {true|false}]`` - Enable NDMP on Vserver::
When the option is set to ``_true_`` , the NDMP daemon handles requests,
and when set to ``_false_`` , the NDMP daemon does not handle requests.
```

```
Enabling and disabling the option is equivalent to executing the following
commands: ``_vserver services ndmp on_`` and ``_vserver services ndmp
off_`` respectively. This option is persistent across reboots. The default
value of this option is ``_false_`` .
+
The `-option-value` for this parameter is either true or false.

``[-preferred-interface-role {undef|cluster|data|node-
mgmt|intercluster|cluster-mgmt}]`` - Preferred Interface Role::
This option allows the user to specify the preferred Logical Interface
(LIF) role while establishing an NDMP data connection channel. The NDMP
data server or the NDMP mover establishes a data channel from the node
that owns the volume or the tape device respectively. This option is used
on the node that owns the volume or the tape device. The order of IP
addresses that are used to establish the data connection depends on the
order of LIF roles specified in this option.
+
The default value for this option for the admin Vserver is
``_intercluster, cluster-mgmt, node-mgmt_``
+
The default value for this option for a data Vserver is ``_intercluster,
data_`` .

``[-secondary-debug-filter <text>]`` - Secondary Debug Filter [.advaced-
priv]#(privilege: advanced)#::
This option allows control on NDMP debug logging. This option takes a
comma separated tag=value pairs. The supported tag is ``_IPADDR_`` which
can be used to specify Vserver IP addresses for which NDMP debugging is
required. If this option is set and the option `debug-enable` is set to
true, then the debug-filter option is applicable to sessions whose control
connection IP addresses match the IP addresses that are listed in the
option. If this option is not set, the debug filter is applicable to all
Vserver sessions. By default, this option does not have a value set.

``[-is-secure-control-connection-enabled {true|false}]`` - Is Secure
Control Connection Enabled::
This option enables NDMP service to accept control connections over secure
sockets on TCP port 30000. This option is persistent across reboots. The
default value of this option is ``_false_`` .

== Examples

The following example show how to enable NDMP on a Vserver and set
authorization type to plaintext :
```

cluster1::> vserver services ndmp modify -vserver vs1 -enable true -authtype plaintext

cluster1::>

```
[[ID57d24740dee35ae9438daacad939067c]]
= vserver services ndmp off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable NDMP service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command is used to disable NDMP service on a specific Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the name of the Vserver.

== Examples

The following example disables NDMP on a specific Vserver:
```

cluster1::> vserver services ndmp off -vserver vs1

```
[[ID38df9c29081608eef26cda79f84585b1]]
= vserver services ndmp on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable NDMP service

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command is used to enable NDMP service on a specific Vserver.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the name of the Vserver.

== Examples

The following example enables NDMP service on a specific Vserver:
```

cluster1::> vserver services ndmp on -vserver vs1

```
[[ID3005f8c352e2faea4abc8f075d22c03e]]
= vserver services ndmp probe
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of NDMP sessions

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The xref:{relative_path}system-services-ndmp-probe.html[system services
ndmp probe] command displays diagnostic information about NDMP sessions
```

```
belonging to a specific Vserver in the cluster. The following fields are
displayed for each of the sessions:

* Vserver
* Session identifier
* NDMP version
* Session authorized
* Data state
* Data operation
* Data server halt reason
* Data server connect type
* Data server connect address
* Data server connect port
* Data bytes processed
* Mover state
* Mover mode
* Mover pause reason
* Mover halt reason
* Mover record size
* Mover record number
* Mover bytes moved
* Mover seek position
* Mover bytes left to read
* Mover window offset
* Mover window length
* Mover position
* Mover SetRecordSize flag
* Mover SetWindow flag
* Mover connect type
* Mover connect address
* Mover connect port
* Effective host
* NDMP client address
* NDMP client port
* SCSI device ID
* SCSI hostadapter
* SCSI target ID
* SCSI LUN ID
* Tape device
* Tape mode
* Node
* Is Secure Control Connection
* Data Backup Mode
* Data Path
* NDMP Source Address
```

```
== Parameters

``[-vserver <vserver name>]`` - Vserver::
This parameter Specifies the Vserver context in which NDMP sessions are
running.

``[-session-id <text>]`` - Session Identifier::
If this parameter is specified, the command displays information about a
specific NDMP session. A session-id is a string used to identify a
particular NDMP session.

``[-ndmp-version <integer>]`` - NDMP Version::
This parameter refers to the NDMP protocol version being used in the
session.

``[-session-authorized {true|false}]`` - Session Authorized::
This parameter indicates whether an NDMP session is authenticated or not.

``[-data-state <component state>]`` - Data State::
This parameter identifies the current state of the data server's state
machine.

``[-data-operation <data operation>]`` - Data Operation::
This parameter identifies the data server's current operation.

``[-data-halt-reason <halt reason>]`` - Data Server Halt Reason::
This parameter identifies the event that caused the data server state
machine to enter the HALTED state.

``[-data-con-addr-type <address type>]`` - Data Server Connect Type::
This parameter specifies the type of data connection established by the
data server. The data connection can be established locally within a given
system or between remote networked systems.

``[-data-con-addr <text>]`` - Data Server Connect Address::
This parameter specifies the connection endpoint information for the data
server's data connection.

``[-data-con-port <integer>]`` - Data Server Connect Port::
This parameter specifies the TCP/IP port that the data server will use
when establishing a data connection.

``[-data-bytes-processed <integer>]`` - Data Bytes Processed::
This parameter represents the cumulative number of data stream bytes
transferred between the backup or recovery method and the data connection
```

```
during the current data operation.

``[-mover-state <component state>]`` - Mover State::
This parameter identifies the current state of the NDMP tape server's
mover state machine.

``[-mover-mode <mover mode>]`` - Mover Mode::
This parameter identifies the direction of the mover data transfer.

``[-mover-pause-reason <pause reason>]`` - Mover Pause Reason::
This parameter identifies the event that caused the mover state machine to
enter the PAUSED state.

``[-mover-halt-reason <halt reason>]`` - Mover Halt Reason::
This parameter field identifies the event that caused the mover state
machine to enter the HALTED state.

``[-mover-record-size <integer>]`` - Mover Record Size::
This parameter represents the current mover record size in bytes.

``[-mover-record-num <integer>]`` - Mover Record Number::
This parameter represents the last tape record processed by the mover.

``[-mover-bytes-moved <integer>]`` - Mover Bytes Moved::
This parameter represents the cumulative number of data stream bytes
written to the data connection or the number of data stream bytes read
from the data connection and written to the tape subsystem, depending on
the mode of mover operation.

``[-mover-seek-position <integer>]`` - Mover Seek Position::
This parameter represents the data stream offset of the first byte the DMA
requested the mover to transfer to the data connection during a mover read
operation.

``[-mover-bytes-left-to-read <integer>]`` - Mover Bytes Left to Read::
This parameter represents the number of data bytes remaining to be
transferred to the data connection to satisfy the current NDMP_MOVER_READ
request.

``[-mover-window-offset <integer>]`` - Mover Window Offset::
This parameter represents the absolute offset of the first byte of the
mover window within the overall data stream.

``[-mover-window-length <integer>]`` - Mover Window Length::
This parameter represents the length of the current mover window in bytes.

``[-mover-position <integer>]`` - Mover Position::
```

```
This parameter can be used to list only those sessions, whose mover
position matches a specific value. Mover-position should be an integer.

``[-mover-setrecordsize-flag {true|false}]`` - Mover SetRecordSize Flag::
This parameter is used by the DMA to establish the record size used for
mover-initiated tape read and write operations.

``[-mover-setwindow-flag {true|false}]`` - Mover SetWindow Flag::
This flag represents whether a mover window has been set or not. A mover
window represents the portion of the overall backup stream that is
accessible to the mover without intervening DMA tape manipulation.

``[-mover-con-addr-type <address type>]`` - Mover Connect Type::
This parameter specifies the type of data connection established by the
mover. The data connection can be established locally within a given
system or between remote networked systems.

``[-mover-con-addr <text>]`` - Mover Connect Address::
This parameter specifies the endpoint address or addresses that the mover
will use when establishing a data connection.

``[-mover-con-port <integer>]`` - Mover Connect Port::
This parameter specifies the TCP/IP port that the mover will use when
establishing a data connection.

``[-eff-host <host type>]`` - Effective Host::
This parameter indicates the host context in which the NDMP session runs.
The valid values are: PRIMARY or PARTNER.

``[-client-addr <text>]`` - NDMP Client Address::
This parameter specifies the client's IP address.

``[-client-port <integer>]`` - NDMP Client Port::
This parameter specifies the client's port number.

``[-spt-device-id <text>]`` - SCSI Device ID::
This parameter specifies the SCSI device ID.

``[-spt-ha <integer>]`` - SCSI Host Adapter::
This parameter specifies the SCSI host adapter.

``[-spt-scsi-id <integer>]`` - SCSI Target ID::
This parameter specifies the SCSI target.

``[-spt-scsi-lun <integer>]`` - SCSI LUN ID::
This parameter specifies the SCSI LUN ID.
```

```
``[-tape-device <text>]`` - Tape Device::
This parameter specifies the name to identify the tape device.

``[-tape-mode <mover mode>]`` - Tape Mode::
This parameter specifies the mode in which tapes are opened.

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays information about the
sessions running on the specified node only. Node should be a valid node
name.

``[-is-secure-control-connection {true|false}]`` - Is Secure Control
Connection::
This parameter specifies whether the control connection is secure or not.

``[-data-backup-mode <text>]`` - Data Backup Mode::
This parameter specifies whether the mode of data backup is Dump or
SMTape.

``[-data-path <text>]`` - Data Path::
This parameter specifies the path of data being backed up.

``[-source-addr <text>]`` - NDMP Source Address::
This parameter specifies the control connection IP address of the NDMP
session.

== Examples

The following example displays diagnostic information about all the
sessions in the cluster:
```

cluster1::> vserver services ndmp probe
Vserver Name: vserver1
Session Identifier: 1000:7445
NDMP Version: 4
Session Authorized: true
Data State: IDLE
Data Operation: NOACTION
Data Server Halt Reason: NA
….
…
Vserver Name: vserver2
Session Identifier: 1000:7446
NDMP Version: 4
Session Authorized: true
Data State: IDLE
Data Operation: NOACTION
Data Server Halt Reason: NA

....
...

```
The following example displays diagnostic information of sessions
associated with Vserver vserver1 only:
```

cluster1::> vserver services ndmp probe -vserver vserver1
Vserver Name: vserver1
Session Identifier: 1000:7445
NDMP Version: 4
Session Authorized: true
Data State: IDLE
Data Operation: NOACTION
Data Server Halt Reason: NA
....
...
....
...

```
== Related Links

* xref:{relative_path}system-services-ndmp-probe.html[system services ndmp
probe]




[[IDf867067faa191b23d7b8bb3a71ec01f4]]
= vserver services ndmp show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NDMP Properties

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command is used to display NDMP options on Vservers.

A combination of parameters can be optionally specified so as to list only
a subset of Vservers where specific values of NDMP options are met. A
short description of each of the options is provided in the parameters
```

```
section.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
If this parameter is specified, the command displays NDMP options for that
Vserver alone.

``[-maxversion <integer>]`` - NDMP Version::
If this parameter is specified, the command displays NDMP options for
Vservers where the highest NDMP protocol version supported matches the
specified input value. The only supported value is 4.

``[-ignore-ctime-enabled {true|false}]`` - Ignore Ctime::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `ignore-ctime-enabled` matches the specified
input value.
+
This option, when ``_true_`` , allows user to exclude files with ctime
changed from storage system' incremental dumps since other processes like
virus scanning often alter the ctime of files. When this option is
``_false_`` , backup on the Vserver will include all files with a change
or modified time later then the last dump in the previous level dump. The
default value is ``_false_`` . This option is persistent across reboots.
+
Most WIN32 APIs are often unaware of the "last changed time", ctime, they
often incorrectly set a later time for files, causing these files to be
included in the Vserver's incremental dumps, making the incremental dump
very large. This is partially defying the purpose of having incremental
dumps, since one uses incremental dumps to speed up the backup by only
dumping files that were truly changed since the last backup.
+
The possible value for this parameter is either true or false.

``[-offset-map-enable {true|false}]`` - Enable Offset Map::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `offset-map-enable` matches the specified
```

```
input value.
+
This option is used to enable or disable generation of the inode offset
map during NDMP based dump backups. The offset map is required to perform
Enhanced Direct Access Restore (DAR) on the backup data. Enhanced DAR
provides support for directory DAR and DAR of files with NT streams. The
default value for this option is ``_true_`` . This option is persistent
across reboots.
+
The possible value for this parameter is either true or false.

``[-tcpnodelay {true|false}]`` - Enable TCP Nodelay::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `tcpnodelay` matches the specified input
value.
+
This parameter Enables/Disables the TCPNODELAY configuration parameter for
the socket between the Vserver and the DMA. When set to true, the Nagle
algorithm is disabled and small packets are sent immediately rather than
held and bundled with other small packets. This optimizes the system for
response time rather than throughput.
+
This option becomes active when the next NDMP session starts. Existing
sessions are unaffected. The default value for this option is ``_false_``
. This option is persistent across reboots.
+
The possible value for this parameter is either true or false.

``[-tcpwinsize <integer>]`` - TCP Window Size::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `tcpwinsize` matches the specified input
value.
+
This option shows the TCP buffer size of the NDMP data connection.The
minimum and maximum values are 8192(8K) and 262,144(256K), respectively.
The default value for this option is 32768(32K).
+
This option is persistent across reboots.
+
The possible value for this parameter is a number between 8192(8K) and
262,144(256K).

``[-data-port-range <text>]`` - Data Port Range::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `data-port-range` matches the specified
input value.
```

```
+
This option shows the port range on which the NDMP server can listen for
data connections.
+
The format of this option is ``_start_port - end_port_`` start_port,
end_port can have values between [1024-65535]; start_port must be lesser
than or equal to end_port. If a valid range is specified, NDMP uses a port
within that range to listen for data connections. A listen request fails
if no ports in the specified range are free.
+
This option is modifiable only from the admin Vserver context and the said
option is applicable for all the data Vservers and the admin Vserver. For
example, if the value of the above option is set with 2000-3000, the same
value will be applicable throughout the cluster. The value ``_all_``
implies that any available port can be used to listen for data
connections. The default value for this option is ``_all_`` . This option
is persistent across reboots.
+
The value for this option is displayed in the format {<start_port>-<end
port> | all }- where start_port, end_port can have values between [1024-
65535]; start_port must be lesser than or equal to end_port.

``[-backup-log-enable {true|false}]`` - Enable Backup Log::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `>backup-log-enable` matches the specified
input value.
+
Backup logging captures important events during dump/restore and records
them in /mroot/etc/log/backup on the root volume. The default value for
this option is ``_true_`` . This option is persistent across reboots.
+
The possible value for this parameter is true/false.

``[-per-qtree-exclude-enable {true|false}]`` - Enable per Qtree
Exclusion::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `per-qtree-exclude-enable` matches the
specified input value.
+
If this option is ``_true_`` , users can specify exclude list on a per
qtree basis to be excluded from backup. This exclude list will override
any values already present due to 'EXCLUDE' environment variable . The
user can specify the exclusion list through a .exclude_list file which
resides at the root of the qtree. The exclusion list can be a list of
files or files that match a specified pattern. The default value for this
option is ``_false_`` . This option is persistent across reboots.
```

```
+
The possible value for this parameter is either true or false.

``[-authtype <NDMP Authentication types>,...]`` - Authentication Type::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `authtype` matches the specified input
value.
+
Allows the administrator to choose the authentication method. NDMP
supports three authentication types: challenge, plaintext and
plaintext_sso. The plaintext_sso authentication type is mutually exclusive
with the other authentication types. By setting the authentication type as
plaintext_sso, the actual password for the user can be used to
authenticate instead of having to generate an NDMP specific password. The
default of this option is ``_challenge_`` . This option is persistent
across reboots.
+
The possible value for this parameter can be {challenge | plaintext |
plaintext_sso | challenge, plaintext | plaintext, challenge}.

``[-debug-enable {true|false}]`` - Enable Debug [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `debug-enable` matches the specified input
value.
+
This option enables debug logging for NDMP. Debug messages will be logged
to the ndmpd log file /mroot/etc/log/mlog/ndmpd.log . The default value
for this option is ``_false_`` .This option is persistent across reboots.
+
The possible value for this parameter is either true or false.

``[-debug-filter <text>]`` - Debug Filter [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `debug-filter` matches the specified input
value.
+
This option controls the NDMP modules for which debug logging is to be
enabled. option-value can take five values for this option : all, none,
normal, backend or "filter-expression".
+
``_all_`` enables debug logging for all modules.
+
``_none_`` disables debug logging for all modules. It is equivalent to
``_modify -vserver vserver_name -debug-enable false_`` .
```

```
+
``_normal_`` is a shortcut option that enables debug logging for all
modules except verbose and io_loop. The equivalent filter string is all-
verbose-io_loop.
+
``_backend_`` is a short cut option that enables debug logging for all
modules except verbose, io_loop, ndmps and ndmpd. The equivalent filter
string is all-verbose-io_loop-ndmps-ndmpp.
+
``_(filter-expression)_`` is a combination of one or more modules for
which debug logs needs to be enabled. Multiple module names can be
combined using following operators :
+
  * - to remove the given module from the list of specified modules in the
filter string. For example the filter all-ndmpp will enable debug logging
for all modules but not ndmpp.
  * ^ to add the given module or modules to the list of modules specified
in the filter string. For example the filter ndmpp^mover^data will enable
debug logging for ndmpp, mover and data.

+
The possible module names and a brief description is given below:-
```

```
+--------------+-------------------------------------+
| Modules      | Description                         |
+--------------+-------------------------------------+
| verbose      | verbose message                     |
| io           | I/O process loop                    |
| io_loop      | I/O process loop verbose messages   |
| ndmps        | NDMP service                        |
| ndmpp        | NDMP Protocol                       |
| rpc          | General RPC service                 |
| fdc_rpc      | RPC to FC driver service            |
| auth         | Authentication                      |
| mover        | NDMP MOVER (tape I/O)               |
| data         | NDMP DATA (backup/restore)          |
| scsi         | NDMP SCSI (robot/tape ops)          |
| bkup_rpc     | RPC to Backup service client        |
| bkup_rpc_s   | RPC to Backup service server        |
| conf         | Debug configure/reconfigure         |
| dblade       | Dblade specific messages            |
| timer        | NDMP server timeout messages        |
| vldb         | VLDB service                        |
| smf          | SMF Gateway messages                |
| common       | NDMP common state                   |
| ext          | NDMP extensions messages            |
| ndmprpc      | NDMP Mhost RPC server               |
+--------------+-------------------------------------+
```

```
+
The default value for this option is ``_none_`` . This option is
persistent across reboots.
+
The possible value for this parameter can be {all | none | normal |
backend |'filter-expression'}.

``[-dump-logical-find <text>]`` - Enable Logical Find for Dump [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `dump-logical-find` matches the specified
input value.
+
This option specifies whether to follow inode-file walk or tree walk for
phase I of the dump. Choosing inode-file walk or tree walk affects the
performance of the dump. This option can take following values:
+
If ``_default_`` is specified, then level 0 and incremental volume as well
```

```
as qtree dumps will use inode walk. All the subtree dumps will use tree
walk.
+
If ``_always_`` is specified, all dumps will follow treewalk.
+
A ``_comma-separated_`` list of values in any combination from the
following list:
+
  * vol_baseline: Level 0 full volume backup will follow treewalk.
  * vol_incr: Incremental full volume backup will follow treewalk.
  * qtree_baseline: Level 0 qtree backup will follow treewalk.
  * qtree_incr: Incremental qtree backup will follow treewalk.



+
The default value for this option is ``_default_`` . This option is
persistent across reboots.
+
The possible value for this parameter could be {default | always |
'vol_baseline' | 'vol_baseline,qtree_baseline' | ...}.

``[-abort-on-disk-error {true|false}]`` - Enable Abort on Disk Error
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `abort-on-disk-error` matches the specified
input value.
+
If this option is ``_true_`` , dump will abort the backup operation on
detection of irrecoverable data blocks in user files. If this option is
``_false_`` , dump will proceed with backup operation - even if
irrecoverable data blocks in user files are detected. On detection of
irrecoverable data blocks, dump will send a log message to DMA and also
log an entry in /mroot/etc/log/backup file. The default value for this
option is ``_false_`` . This option is persistent across reboots.
+
The value for this parameter is either true or false.

``[-fh-dir-retry-interval <integer>]`` - FH Throttle Value for Dir
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `fh-dir-retry-interval` matches the
specified input value.
+
NDMP protocol sends back file history information for all directories in
phase 3 of dump to DMA. In the presence of slow DMA or high latency
networks, the amount of file history being generated exceeds the amount
```

```
being consumed by the DMA. To handle a slow reader, a flow control
mechanism is now introduced where file history generation is throttled
when a DMA is slow in consuming them. The value for this option indicates
how frequently should the file history be resent if it was throttled. The
default value is 250 milliseconds. This option is persistent across
reboots.
+
The value for this parameter is a number.

``[-fh-node-retry-interval <integer>]`` - FH Throttle Value for Node
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `fh-node-retry-interval` matches the
specified input value.
+
NDMP protocol sends back file history information for all files in phase 4
of dump to DMA. In the presence of slow DMA or high latency networks, the
amount of file history being generated exceeds the amount being consumed
by the DMA. To handle slow reader conditions, a flow control mechanism is
now introduced where file history generation is throttled when a DMA is
slow in consuming them. The value for this option indicates how frequently
should the file history be resent if it was throttled. The default value
is 250 milliseconds. This option is persistent across reboots.
+
The value for this parameter is a number.

``[-restore-vm-cache-size <integer>]`` - Restore VM File Cache Size
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `restore-vm-cache-size` matches the
specified input value.
+
This option mandates the number of WAFL buffers pinned in memory by
various meta-files used by logical restore. The minimum and maximum values
are 4 and 1024, respectively. The default value for this option is 64.
This option is persistent across reboots.
+
Depending on the value of this option, various meta-files are assigned a
number of WAFL buffers that need to be pinned in memory.
```

```
+-----------------+---------------------------------+
| Meta-filename   |     Number of WAFL buffers to   |
|                 |      be pinned in memory        |
+-----------------+---------------------------------+
| dumpmap         | ndmpd.restore.vm_cache_size     |
| filemap         | ndmpd.restore.vm_cache_size     |
| aclfile_map     | ndmpd.restore.vm_cache_size     |
| inomap          | ndmpd.restore.vm_cache_size / 2 |
| basemap         | ndmpd.restore.vm_cache_size / 2 |
| flipmap         | ndmpd.restore.vm_cache_size / 2 |
| revmap          | ndmpd.restore.vm_cache_size / 2 |
| clrimap         | ndmpd.restore.vm_cache_size / 4 |
| mfp_for_inotab  | ndmpd.restore.vm_cache_size / 4 |
| map             | ndmpd.restore.vm_cache_size / 4 |
| offsetfile_map  | ndmpd.restore.vm_cache_size / 4 |
+-----------------+---------------------------------+
```

```
+
The possible value for this parameter is a number between 4 and 1024.

``[-enable {true|false}]`` - Enable NDMP on Vserver::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `enable` matches the specified input value.
+
When the option is set to ``_true_`` , the NDMP daemon handles requests,
and when set to ``_false_`` , the NDMP daemon does not handle requests.
Enabling and disabling the option is equivalent to executing the following
commands: ``_vserver services ndmp on_`` and ``_vserver services ndmp
off_`` respectively. This option is persistent across reboots. The default
value of this option is ``_false_`` .
+
The value for this parameter is either true or false.

``[-preferred-interface-role {undef|cluster|data|node-
mgmt|intercluster|cluster-mgmt}]`` - Preferred Interface Role::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `preferred-interface-role` matches the
specified input value.
+
This option allows the user to specify the preferred Logical Interface
(LIF) role while establishing an NDMP data connection channel. The NDMP
data server or the NDMP mover establishes a data channel from the node
that owns the volume or the tape device respectively. This option is used
on the node that owns the volume or the tape device. The order of IP
```

```
addresses that are used to establish the data connection depends on the
order of LIF roles specified in this option.
+
The default value for this option for the admin Vserver is
``_intercluster, cluster-mgmt, node-mgmt_``
+
The default value for this option for a data Vserver is ``_intercluster,
data_`` .

``[-secondary-debug-filter <text>]`` - Secondary Debug Filter [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `secondary-debug-filter` matches the
specified input value.
+
This option allows control on NDMP debug logging. This option takes a
comma separated tag=value pairs. The supported tag is ``_IPADDR_`` which
can be used to specify Vserver IP addresses for which NDMP debugging is
required. If this option is set and the option `debug-enable` is set to
true, then the debug-filter option is applicable to sessions whose control
connection IP addresses match the IP addresses that are listed in the
option. If this option is not set, the debug filter is applicable to all
Vserver sessions. By default, this option does not have a value set.

``[-is-secure-control-connection-enabled {true|false}]`` - Is Secure
Control Connection Enabled::
If this parameter is specified, the command displays NDMP options for
Vservers, where the value for `is-secure-control-connection-enabled`
matches the specified input value.
+
This option enables NDMP service to accept control connections over secure
sockets on TCP port 30000. This option is persistent across reboots. The
default value of this option is ``_false_`` .

== Examples

The following example displays NDMP options for the Vserver(s).
```

cluster1::> vserver services ndmp show

VServer Enabled Authentication type
--- --- ---
cluster true plaintext
vs1 true plaintext
vs2 true plaintext
3 entries were displayed.

cluster1::>

```
The following example displays detailed NDMP options for a Vserver.
```

cluster1::*> vserver services ndmp show -vserver vs1 -instance
Vserver: vs1
NDMP Version: 4
Ignore Ctime: false
Enable Offset Map: true
Enable TCP Nodelay: false
TCP Window Size: 32768
Data Port Range: all
Enable Backup Log: true
Enable per Qtree Exclusion: false
Authentication Type: plaintext
Enable Debug: false
Debug Filter: none
Enable Logical Find for Dump: default
Enable Abort on Disk Error: false
FH Throttle Value for Dir: 250
FH Throttle Value for Node: 250
Restore VM File Cache Size: 64
Enable Logging of VM Stats for Dump: false
Enable NDMP on Vserver: true
Preferred Interface Role: intercluster, data
Secondary Debug Filter: -
Is Secure Control Connection Enabled: false
cluster1::*>

```
[[IDd1bc5054fddcbfa9455a10f0fb7ecd62]]
= vserver services ndmp status
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of NDMP sessions

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services ndmp status` command lists NDMP sessions belonging
to a specific Vserver in the cluster. By default it lists the following
details about the active sessions:

* Vserver Name
```

```
* Session ID

A combination of parameters can be optionally supplied so as to list only
those sessions which match specific conditions. A short description of
each of the parameter is provided in the parameters section.

== Parameters

{ ``[-fields <fieldname>,...]``::
This optional parameter specifies which all additional fields to display.
Any combination of the following fields are valid:
+
  * ndmp-version
  * session-authorized
  * data-state
  * data-operation
  * data-halt-reason
  * data-con-addr-type
  * data-con-addr
  * data-con-port
  * data-bytes-processed
  * mover-state
  * mover-mode
  * mover-pause-reason
  * mover-halt-reason
  * mover-record-size
  * mover-record-num
  * mover-bytes-moved
  * mover-seek-position
  * mover-bytes-left-to-read
  * mover-window-offset
  * mover-window-length
  * mover-position
  * mover-setrecordsize-flag
  * mover-setwindow-flag
  * mover-con-addr-type
  * mover-con-addr
  * mover-con-port
  * eff-host
  * client-addr
  * client-port
  * spt-device-id
  * spt-ha
  * spt-scsi-id
  * spt-scsi-lun
  * tape-device
```

```
  * tape-modes
  * node
  * is-secure-control-connection
  * data-backup-mode
  * data-path
  * source-addr



| ``[-instance  ]`` }::
If this parameter is specified, the command displays detailed information
about all the active sessions.

``[-vserver <vserver name>]`` - Vserver::
Specifies the Vserver context in which NDMP sessions are running.

``[-session-id <text>]`` - Session Identifier::
If this parameter is specified, the command displays information about
specific NDMP session. A session-id is a string used to identify a
particular NDMP session.

``[-ndmp-version <integer>]`` - NDMP Version::
This parameter refers to the NDMP protocol version being used in the
session.

``[-session-authorized {true|false}]`` - Session Authorized::
This field indicates whether an NDMP session is authenticated or not.

``[-data-state <component state>]`` - Data State::
This field identifies the current state of the data server's state
machine.

``[-data-operation <data operation>]`` - Data Operation::
This field identifies the data server's current operation.

``[-data-halt-reason <halt reason>]`` - Data Server Halt Reason::
This field identifies the event that caused the data server state machine
to enter the HALTED state.

``[-data-con-addr-type <address type>]`` - Data Server Connect Type::
This field specifies the type of data connection established by the data
server. The data connection can be established locally within a given
system or between remote networked systems.

``[-data-con-addr <text>]`` - Data Server Connect Address::
This specifies the connection endpoint information for the data server's
data connection.
```

```
``[-data-con-port <integer>]`` - Data Server Connect Port::
This specifies the TCP/IP port that the data server will use when
establishing a data connection.

``[-data-bytes-processed <integer>]`` - Data Bytes Processed::
This field represents the cumulative number of data stream bytes
transferred between the backup or recovery method and the data connection
during the current data operation.

``[-mover-state <component state>]`` - Mover State::
This parameter identifies the current state of the NDMP tape server's
mover state machine.

``[-mover-mode <mover mode>]`` - Mover Mode::
This parameter identifies the direction of the mover data transfer.

``[-mover-pause-reason <pause reason>]`` - Mover Pause Reason::
This parameter identifies the event that caused the mover state machine to
enter the PAUSED state.

``[-mover-halt-reason <halt reason>]`` - Mover Halt Reason::
This integer field identifies the event that caused the mover state
machine to enter the HALTED state.

``[-mover-record-size <integer>]`` - Mover Record Size::
This field represents the current mover record size in bytes.

``[-mover-record-num <integer>]`` - Mover Record Number::
This field represents the last tape record processed by the mover.

``[-mover-bytes-moved <integer>]`` - Mover Bytes Moved::
This field represents the cumulative number of data stream bytes written
to the data connection or the number of data stream bytes read from the
data connection and written to the tape subsystem, depending on the mode
of mover operation.

``[-mover-seek-position <integer>]`` - Mover Seek Position::
This field represents the data stream offset of the first byte the DMA
requested the mover to transfer to the data connection during a mover read
operation.

``[-mover-bytes-left-to-read <integer>]`` - Mover Bytes Left to Read::
This field represents the number of data bytes remaining to be transferred
to the data connection to satisfy the current NDMP_MOVER_READ request.

``[-mover-window-offset <integer>]`` - Mover Window Offset::
```

```
This field represents the absolute offset of the first byte of the mover
window within the overall data stream.

``[-mover-window-length <integer>]`` - Mover Window Length::
This field represents the length of the current mover window in bytes.

``[-mover-position <integer>]`` - Mover Position::
This parameter can be used to list only those sessions, whose mover
position matches a specific value. Mover-position should be an integer.

``[-mover-setrecordsize-flag {true|false}]`` - Mover SetRecordSize Flag::
This field is used by the DMA to establish the record size used for mover-
initiated tape read and write operations.

``[-mover-setwindow-flag {true|false}]`` - Mover SetWindow Flag::
This flag represents whether a mover window has been set or not. A mover
window represents the portion of the overall backup stream that is
accessible to the mover without intervening DMA tape manipulation.

``[-mover-con-addr-type <address type>]`` - Mover Connect Type::
This field specifies the type of data connection established by the mover.
The data connection can be established locally within a given system or
between remote networked systems.

``[-mover-con-addr <text>]`` - Mover Connect Address::
This specifies the endpoint address or addresses that the mover will use
when establishing a data connection.

``[-mover-con-port <integer>]`` - Mover Connect Port::
This specifies the TCP/IP port that the mover will use when establishing a
data connection.

``[-eff-host <host type>]`` - Effective Host::
This field indicates the host context in which the NDMP session runs. The
valid values are: PRIMARY or PARTNER.

``[-client-addr <text>]`` - NDMP Client Address::
This parameter specifies the client's IP address.

``[-client-port <integer>]`` - NDMP Client Port::
This parameter specifies the client's port number.

``[-spt-device-id <text>]`` - SCSI Device ID::
This parameter specifies the SCSI device ID.

``[-spt-ha <integer>]`` - SCSI Host Adapter::
This parameter specifies the SCSI host adapter.
```

```
``[-spt-scsi-id <integer>]`` - SCSI Target ID::
This parameter specifies the SCSI target.

``[-spt-scsi-lun <integer>]`` - SCSI LUN ID::
This parameter specifies the SCSI LUN ID.

``[-tape-device <text>]`` - Tape Device::
This parameter specifies the name to identify the tape device.

``[-tape-mode <mover mode>]`` - Tape Mode::
This parameter specifies the mode in which tapes are opened.

``[-node {<nodename>|local}]`` - Node::
If this parameter is specified, the command displays information about the
sessions running on the specified node only. Node should be a valid node
name.

``[-is-secure-control-connection {true|false}]`` - Is Secure Control
Connection::
This parameter specifies whether the control connection is secure or not.

``[-data-backup-mode <text>]`` - Data Backup Mode::
This parameter specifies whether the mode of data backup is Dump or
SMTape.

``[-data-path <text>]`` - Data Path::
This parameter specifies the path of data being backed up.

``[-source-addr <text>]`` - NDMP Source Address::
This parameter specifies the control connection IP address of the NDMP
session.

== Examples

The following example displays all the NDMP sessions on the cluster:
```

cluster1::> vserver services ndmp status
Session
Vserver Id
--- ---
vserver1 1000:7445
vserver2 1000:7446
vserver2 1000:7447
3 entries were displayed.

```
The following example shows how to display only the sessions running
belonging to Vserver vserver2:
```

cluster1::> vserver services ndmp status -vserver vserver2
Session
Vserver Id
---------- ----------
vserver2 1000:7446
vserver2 1000:7447
2 entries were displayed.

```
[[ID4e06c9eff672d117e2ebff7aaa2a3974]]
= vserver services ndmp extensions modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify NDMP extension status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command is used to enable/disable an NDMP extension in the Vserver-
aware NDMP mode.

== Parameters

``[-is-extension-0x2050-enabled {true|false}]`` - Is Extension 0x2050
Enabled [.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command can be used to modify the
status of the extension in the Vserver-aware mode.

== Examples

The following example shows how to enable NDMP extension 0x2050 in the
Vserver-aware NDMP mode of operation:
```

cluster1::> vserver services ndmp extension modify -is-extension-0x2050-enabled true
cluster1::>

```
[[IDb72f6dc2b002af765967bef90c56fc7d]]
= vserver services ndmp extensions show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NDMP extension status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command displays whether an NDMP extension is enabled in the Vserver-
aware NDMP mode or not.

== Examples

The following example shows how to check the status of NDMP extension
0x2050 in a cluster :
```

cluster1::> vserver services ndmp extension show
Is Extension 0x2050 Enabled: true
cluster1::>

```
[[ID52cc3e4bf4c43b0ab19235f33335efc5]]
= vserver services ndmp log start
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start logging for the specified NDMP session

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command is used to start logging on an active NDMP session on a
vserver.
```

```
== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the Vserver.

``-session-id <text>`` - Session Identifier [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the NDMP session-id on which logging needs to be
started.

``-filter <text>`` - Level Filter [.advaced-priv]#(privilege: advanced)#::
Use this parameter to specify the filter for a particular session ID. This
parameter controls the NDMP modules for which logging is to be enabled.
This parameter can take five values. They are as follow : ``_all_`` ,
``_none_`` , ``_normal_`` , ``_backend_`` or ``_"filter-expression"_`` .
The default value for this is ``_none_`` .
+
  * ``_all_`` turns on logging for all modules.
  * ``_none_`` disables logging for all modules.
  * ``_normal_`` is a short cut parameter that enables logging for all
modules except ``_verbose_`` and ``_io_loop_`` . The equivalent filter
string is ``_all-verbose-io_loop_``
  * ``_backend_`` is a short cut parameter that enables logging for all
modules except ``_verbose_`` , ``_io_loop_`` , ``_ndmps_`` and ``_ndmpd_``
. The equivalent filter string is ``_all-verbose-io_loop-ndmps-ndmpp_``
  * ``_(filter-expression)_`` is a combination of one or more modules for
which logs needs to be enabled. Multiple module names can be combined
using following operators :
+
    * ``_-_`` to remove the given module from the list of specified
modules in the filter string. For example the filter ``_all-ndmpp_`` will
enable logging for all modules but not ``_ndmpp_`` .
    * ``_^_`` to add the given module or modules to the list of modules
specified in the filter string. For example the filter
``_ndmpp^mover^data_`` will enable logging for ``_ndmpp_`` , ``_mover_``
and ``_data_`` .



+
The possible module names and a brief description is given below:
```

```
+--------------+-------------------------------------+
| Modules      | Description                         |
+--------------+-------------------------------------+
| verbose      | verbose message                     |
| io           | I/O process loop                    |
| io_loop      | I/O process loop verbose messages   |
| ndmps        | NDMP service                        |
| ndmpp        | NDMP Protocol                       |
| rpc          | General RPC service                 |
| fdc_rpc      | RPC to FC driver service            |
| auth         | Authentication                      |
| mover        | NDMP MOVER (tape I/O)               |
| data         | NDMP DATA (backup/restore)          |
| scsi         | NDMP SCSI (robot/tape ops)          |
| bkup_rpc     | RPC to Backup service client        |
| bkup_rpc_s   | RPC to Backup service server        |
| cleaner      | Backup/Mover session cleaner        |
| conf         | Debug configure/reconfigure         |
| dblade       | Dblade specific messages            |
| timer        | NDMP server timeout messages        |
| vldb         | VLDB service                        |
| smf          | SMF Gateway messages                |
| vol          | VOL OPS service                     |
| sv           | SnapVault NDMP extension            |
| common       | NDMP common state                   |
| ext          | NDMP extensions messages            |
| sm           | SnapMirror NDMP extension           |
| ndmprpc      | NDMP Mhost RPC server               |
+--------------+-------------------------------------+
```

```
== Examples

The following example shows how to start logging on a specific NDMP
session 1000:35512, running on vserver cluster1-01 with filter all.
```

cluster1::*> vserver services ndmp log start -vserver cluster1-01 -session-id 1000:35512 -filter all

```
[[ID9da93d8e09b40392478dfd7d95139b79]]
= vserver services ndmp log stop
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop logging for the specified NDMP session

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

This command is used to stop logging on an active NDMP session on a
vserver.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the name of the Vserver.

``-session-id <text>`` - Session Identifier [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the NDMP session-id on which logging needs to be
stopped.

== Examples

The following example shows how to stop logging on a specific NDMP session
1000:35512 , running on vserver cluster1-01.
```

cluster1::*> vserver services ndmp log stop -vserver cluster1-01 -session-id 1000:35512

```
[[ID6d4622279d5dfb931344e9a54f5ceb78]]
= vserver services ndmp restartable-backup delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an NDMP restartable backup context

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services ndmp restartable-backup delete` command deletes an
NDMP restartable backup context. The `-force` flag can be used to forcibly
destroy a NDMP restartable backup context.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
This parameter specifies the name of the Vserver for which NDMP
restartable backup context is to be deleted.

``-context-id <UUID>`` - Context Identifier::
This parameter specifies the NDMP restartable backup context ID which
needs to be deleted.

``[-force <true>]`` - Force Delete [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the context is deleted even if there are
internal errors.

== Examples

The following example shows how to delete an NDMP restartable backup
context:
```

cluster1::> vserver services ndmp restartable-backup delete -vserver cluster1-01 -context-id 0f8f5c44-d540-11e5-8c45-005056963504
cluster1::>

```
[[ID2c1a2bfcb8474424767f2f7f36d6346e]]
= vserver services ndmp restartable-backup show
```

```
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display NDMP restartable backup contexts

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver services ndmp restartable-backup show` command lists the NDMP
restartable backup contexts present in the cluster. By default it lists
the following details about the context:

* Vserver Name
* Context Identifier
* Is Cleanup Pending?

A combination of parameters can be optionally supplied so as to list only
those contexts which match specific conditions. A short description of
each of the parameter is provided in the parameters section.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
If this parameter is specified, the command displays NDMP restartable
backup contexts that match the specified Vserver.

``[-context-id <UUID>]`` - Context Identifier::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for `-context-id` matches the specified
input value.
+
This parameter specifies the UUID of NDMP restartable backup contexts.
```

```
``[-volume <volume name>]`` - Volume Name::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for volume matches the specified input
value.
+
This parameter specifies the volume path information

``[-is-cleanup-pending {true|false}]`` - Is Cleanup Pending?::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for `-is-cleanup-pending` matches the
specified input value.
+
This parameter indicates whether the context is being deleted.

``[-engine-type <text>]`` - Backup Engine Type::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for engine-type matches the specified
input value.
+
This parameter specifies the backup engine type.

``[-auto-snapshot {true|false}]`` - Is Snapshot Copy Auto-created?::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for auto-snapshot matches the specified
input value.
+
This parameter indicates if the snapshot was created by DUMP engine.

``[-no-acls {true|false}]`` - Is NO_ACLS Set? [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for no-acls matches the specified input
value.
+
This parameter specifies if NO_ACLS environment variable is set.

``[-dump-path <text>]`` - Dump Path::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for dumppath matches the specified input
value.
+
This parameter represents the corresponding local volume path which is
being backed up.

``[-backup-level <integer>]`` - Incremental Backup Level ID::
If this parameter is specified, the command displays NDMP restartable
```

```
backup contexts where the value for backup-level matches the specified
input value.
+
This parameter specifies the backup level.

``[-dump-date <integer>]`` - Dump Date [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for dumpdate matches the specified input
value.
+
This parameter specifies the dumpdate value in epoch.

``[-base-date <integer>]`` - Base Date [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for basedate matches the specified input
value.
+
This parameter specifies the basedate value in epoch.

``[-update-dump-dates {true|false}]`` - Dump Dates Require Update?
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for update-dumpdates matches the specified
input value.
+
This parameter indicates if dumpdates needs to be updated.

``[-dump-name <text>]`` - Dump Name::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for dumpname matches the specified input
value.
+
This parameter indicates the name for the dump instance.

``[-all-non-qtree {true|false}]`` - Is NON_QUOTA_QTREE Set? [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for all-non-qtree matches the specified
input value.
+
This parameter indicates if NON_QUOTA_TREE environment variable is set.

``[-print-options <integer>]`` - Backup Log Level [.advaced-
priv]#(privilege: advanced)#::
```

```
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for print-options matches the specified
input value.
+
This parameter specifies the logging level during dump.

``[-last-update <integer>]`` - Context Last Updated Time::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for last-update matches the specified
input value.
+
This parameter specifies the last time(in epoch) when the context was
modified.

``[-has-offset-map {true|false}]`` - Has Offset Map?::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for has-offset-map matches the specified
input value.
+
This parameter indicates if offset map is present in the backup image.

``[-offset-verify {true|false}]`` - Offset Verify::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for offset-verify matches the specified
input value.
+
This parameter indicates if offset map is successfully verified during
backup.

``[-ndmp-env-keys <text>,...]`` - NDMP Environment Keys [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for ndmpenvkeys matches the specified
input value.
+
This parameter represents the list of NDMP environment variables set
during backup.

``[-ndmp-env-values <text>,...]`` - NDMP Environment Values [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for ndmpenvvalues matches the specified
input value.
+
This parameter represents the values set for the NDMP environment
variables.
```

```
``[-ndmp-env-count <integer>]`` - Count of NDMP Environment Variables
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for ndmpenvcount matches the specified
input value.
+
This parameter represents the number of NDMP environment variables set
during backup.

``[-is-restartable {true|false}]`` - Is Context Restartable?::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for is-restartable matches the specified
input value.
+
This parameter indicates if the NDMP restartable backup context is
restartable.

``[-is-busy {true|false}]`` - Is Context Busy?::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for is-busy matches the specified input
value.
+
This parameter indicates if the NDMP restartable backup context is busy.

``[-multi-subtree {true|false}]`` - Is MULTI_SUBTREE_NAMES Set? [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for multi-subtree matches the specified
input value.
+
This parameter indicates if the NDMP environment variable
MULTI_SUBTREE_NAMES is set.

``[-logical-find {true|false}]`` - Is LOGICAL_FIND Set? [.advaced-
priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for logical-find matches the specified
input value.
+
This parameter indicates if the NDMP environment variable LOGICAL_FIND is
set.

``[-exclude-list <text>]`` - Is EXCLUDE Set? [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for exclude-list matches the specified
```

```
input value.
+
This parameter represents the value of the the NDMP environment variable
EXCLUDE.

``[-restart-pass <integer>]`` - Restart Pass::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for restart-pass matches the specified
input value.
+
This parameter specifies the dump phase from which to restart.

``[-backup-results <integer>]`` - Status of Backup::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for backup-results matches the specified
input value.
+
This parameter specifies the status of the backup.

``[-snap-name <text>]`` - Snapshot Copy Name::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for snap-name matches the specified input
value.
+
This parameter specifies the name of the snapshot.

``[-is-dp-vol {true|false}]`` - Is DP Volume? [.advaced-priv]#(privilege:
advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for is-dp-vol matches the specified input
value.
+
This parameter indicates if the volume specified in the NDMP restartable
context is of type DP.

``[-context-status <integer>]`` - State of the Context::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for context-status matches the specified
input value.
+
This parameter specifies the state of the NDMP restartable context.

``[-is-fg-vol {true|false}]`` - Is FlexGroup Volume?::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for is-fg-vol matches the specified input
value.
```

```
+
This parameter indicates if the volume specified in the NDMP restartable
context is a FlexGroup volume or not.

``[-is-64bit-atime-backup {true|false}]`` - Is 64bit Atime Backup?
[.advaced-priv]#(privilege: advanced)#::
If this parameter is specified, the command displays NDMP restartable
backup contexts where the value for is-64bit-atime-backup matches the
specified input value.
+
This parameter indicates if the backup specified in the NDMP restartable
context is a 64bit atime backup or not.

== Examples

The following example displays all the NDMP restartable contexts on the
cluster:
```

cluster1::> vserver services ndmp restartable-backup show
Vserver Context Identifier Is Cleanup Pending?
---------- ---------- ----------
vserver1 53a6760e-d245-11e5-a33b-005056bb2685 false
vserver2 68902360-d245-11e5-a33b-005056bb2685 true
vserver2 d7b74e0d-d24c-11e5-a33b-005056bb2685 false
3 entries were displayed.

```
The following example shows how to display only the contexts belonging to
Vserver vserver2:
```

cluster1::> vserver services ndmp restartable-backup show -vserver vserver2
Vserver Context Identifier Is Cleanup Pending?
---------- ---------- ----------
vserver2 68902360-d245-11e5-a33b-005056bb2685 true
vserver2 d7b74e0d-d24c-11e5-a33b-005056bb2685 false
2 entries were displayed.

```
[[ID0163e9c5242ef02d5458564b07572199]]
= vserver services web modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify the configuration of web services

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command modifies the availability of the web services on Vservers.
Only the services that are installed on every node in the cluster can be
configured on Vservers whose type is not 'node'.  Enabled services must
include authorization configuration in the `vserver services web access`
command for the services to be externally available.

== Parameters

``-vserver <Vserver Name>`` - Vserver::
Identifies a Vserver for hosting a specific web service.

``-name <text>`` - Service Name::
Identifies the name of the web service.

``[-enabled {true|false}]`` - Enabled::
Defines the availability of a service on the Vserver. Disabled services
are not accessible through the Vserver's network interfaces. This
parameter's default value is dependent on the service. In general,
services that provide commonly used features are enabled by default.

``[-ssl-only {true|false}]`` - SSL Only::
Defines the encryption enforcement policy for a service on the Vserver.
Services for which this parameter is set to true support SSL only and
cannot be used over unencrypted HTTP. The default for this value is
'false'.

== Examples

The following command sets access to the web port to SSL only:
```

cluster1::> vserver services web modify -vserver vs1 -name portal -ssl-only true

```
[[IDaae281667db63b4d85e564f1951ae9e3]]
= vserver services web show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the current configuration of web services

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the availability of the web services on Vservers.
Only the services that are installed on every node in the cluster can be
configured on Vservers whose type is not 'node'.  Enabled services must
include authorization configuration in the `vserver services web access`
command for the services to be externally available.

== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Identifies a Vserver for hosting a specific web service.

``[-name <text>]`` - Service Name::
Identifies the name of the web service.

``[-type <vserver type>]`` - Type of Vserver::
Identifies the type of Vserver on which the service is hosted.

``[-version <text>]`` - Version of Web Service::
Defines the version number of the service in the format of
major.minor.patch.
```

```
``[-description <text>]`` - Description of Web Service::
Provides a short description of the web service.

``[-long-description <text>]`` - Long Description of Web Service::
Provides a long description of the web service.

``[-requires <requirement>,...]`` - Service Requirements::
Defines the list of requirements that must be met for the service to be
successfully executed. Requirements are defined as a service name, a
comparison operator (<=>), and a version number.

``[-default-roles <text>,...]`` - Default Authorized Roles::
Defines the roles that are automatically granted access to the service in
the xref:{relative_path}vserver-services-web-access-show.html[vserver
services web access show] configuration.

``[-enabled {true|false}]`` - Enabled::
Defines the availability of a service on the Vserver. Disabled services
are not accessible through the Vserver's network interfaces. This
parameter's default value is dependent on the service. In general,
services that provide commonly used features are enabled by default.

``[-ssl-only {true|false}]`` - SSL Only::
Defines the encryption enforcement policy for a service on the Vserver.
Services for which this parameter is set to true support SSL only and
cannot be used over unencrypted HTTP. The default for this value is
'false'.

== Examples

This example displays the availability of the web services on the
Vservers.
```

cluster1::vserver services web> show
Vserver Type Service Name Description Enabled
--- --- --- --- ---
cluster1 admin cem OBSOLETE true
cluster1 admin ontapi Remote Administrative API true
cluster1 admin portal Data ONTAP Web Services true
Portal
n6070-8 node cem OBSOLETE true
n6070-8 node ontapi Remote Administrative API true
n6070-8 node portal Data ONTAP Web Services true
Portal
n6070-8 node spi Service Processor false
Infrastructure
n6070-8 node supdiag Support Diagnostics true

Support
n6070-9 node cem OBSOLETE true
n6070-9 node ontapi Remote Administrative API true
n6070-9 node portal Data ONTAP Web Services true
Portal
n6070-9 node spi Service Processor false
Infrastructure
n6070-9 node supdiag Support Diagnostics false
Support
13 entries were displayed.

```
== Related Links

* xref:{relative_path}vserver-services-web-access-show.html[vserver
services web access show]




[[ID777d5bf649416d815a2f6c24d561f321]]
= vserver services web access create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Authorize a new role for web service access

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command authorizes roles to access the Vserver's web services. For
the user to access services that require authentication, the user's roles,
as defined by xref:{relative_path}security-login-show.html[security login
show] , must be included in this configuration.

NOTE: Node Vserver services are authorized with the data Vserver's roles.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
Identifies a Vserver for hosting a specific web service.
```

```
``-name <text>`` - Service Name::
Identifies the name of the web service.

``-role <text>`` - Role Name::
Identifies the new role to be authorized for this service.

== Examples

The following example authorizes the role ``_auditor_`` - created
previously - for the web service:
```

cluster1::> vserver services web access create -name ontapi -role auditor

```
== Related Links

* xref:{relative_path}security-login-show.html[security login show]




[[ID5b460fe5e1b682b90ea72a5eaae927da]]
= vserver services web access delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove role authorization for web service access

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command removes the authorization of a role from the Vserver's web
services. A service for which no roles are defined has a single role of
'none' automatically displayed in this configuration.

NOTE: Node Vserver services are authorized with the data Vserver's roles.



== Parameters

``-vserver <Vserver Name>`` - Vserver::
```

```
Identifies a Vserver for hosting a specific web service.

``-name <text>`` - Service Name::
Identifies the name of the web service.

``-role <text>`` - Role Name::
Identifies the role whose authorization is to be removed. You cannot
remove the authorization of the role 'none'. Use
xref:{relative_path}vserver-services-web-access-create.html[vserver
services web access create] to authorize access for the role.

== Examples

The following example removes authorization for the role ``_auditor_`` for
the web service:
```

cluster1::> vserver services web access delete -name ontapi -role auditor

```
== Related Links

* xref:{relative_path}vserver-services-web-access-create.html[vserver
services web access create]




[[IDf8fd832af5ee3090c1fdebfeceb219ea]]
= vserver services web access show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display web service authorization for user roles

*Availability:* This command is available to _cluster_ administrators at
the _admin_ privilege level.

== Description

This command displays the roles that are authorized to access the
Vserver's web services. For the user to access services that require
authentication, the user's roles, as defined by
xref:{relative_path}security-login-show.html[security login show] , must
be included in this configuration.
```

```
NOTE: Node Vserver services are authorized with the data Vserver's roles.



== 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.

``[-vserver <Vserver Name>]`` - Vserver::
Identifies a Vserver for hosting a specific web service.

``[-name <text>]`` - Service Name::
Identifies the name of the web service.

``[-role <text>]`` - Role Name::
Identifies a role assigned for accessing the service. A service without
any authorizations has a role of 'none' assigned to it automatically.

``[-type <vserver type>]`` - Type of Vserver::
Identifies the type of Vserver on which the service is hosted.

== Examples

The following example displays the roles that are authorized to access the
web services.
```

cluster1::vserver services web access> show
Vserver Type Service Name Role
--- --- --- ---
cluster1 admin cem none
cluster1 admin ontapi readonly
cluster1 admin portal none
cluster1 admin spi none
cluster1 admin supdiag none
vs0 cluster ontapi admin
6 entries were displayed.

cluster1::vserver services web access>

```
== Related Links

* xref:{relative_path}security-login-show.html[security login show]




:leveloffset: -1


= vserver smtape commands

:leveloffset: +1


[[ID22a7be35233bf696787fbabbc8315456]]
= vserver smtape break
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Make a restored volume read-write

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

This command breaks the relationship between the tape backup of a volume
and a restored volume, changing the restored volume from read-only to
read/write.

== Parameters

``-vserver <vserver name>`` - Vserver Name::
Use this parameter to specify the Vserver name on which the volume is
located.

``-volume <volume name>`` - Volume Name::
Use this parameter to specify the name of the read-only volume that needs
to be changed into a read/writeable volume after an smtape restore.

== Examples

Make the read-only volume ``_datavol_`` on Vserver ``_vserver0_``
```

```
writeable after a restore.
```

cluster1::> vserver smtape break -vserver vserver0 -volume datavol
[Job 84] Job succeeded: SnapMirror Break Succeeded

```
:leveloffset: -1


= vserver snapdiff-rpc-server commands

:leveloffset: +1


[[IDfb1cf5ab0dfb6925db79d5634af13aec]]
= vserver snapdiff-rpc-server off
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Stop the SnapDiff RPC server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver snapdiff-rpc-server off` command turns the SnapDiff RPC
server off.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver for which you want to turn the
SnapDiff RPC server off.

== Examples

The following example turns the SnapDiff RPC server off for a Vserver
named vs0:
```

cluster1::> vserver snapdiff-rpc-server off -vserver vs0

```
[[IDb9544f3a703640ce9778f47874c038d5]]
= vserver snapdiff-rpc-server on
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Start the SnapDiff RPC Server

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver snapdiff-rpc-server on` command turns the SnapDiff RPC server
on.

== Parameters

``-vserver <vserver name>`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
This parameter specifies the Vserver for which you want to turn the
SnapDiff RPC server on.

== Examples

The following example enables the SnapDiff RPC server access for a Vserver
named vs0:
```

cluster1::> vserver snapdiff-rpc-server on -vserver vs0

```
[[ID20b72b94e946ec6437d085b850c21c6d]]
= vserver snapdiff-rpc-server show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display the SnapDiff RPC server configurations of Vservers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.
```

```
== Description

The `vserver snapdiff-rpc-server show` command displays the state of the
SnapDiff RPC server for all the Vservers. The command output depends on
the parameter or parameters specified with the command. If no parameters
are specified, the command displays the following information about all
the configured Vservers:

* Vserver name
* Whether SnapDiff RPC server access is enabled

You can specify additional parameters to display only the information that
matches those parameters. For instance, to display the information only
for the Vservers that have access enabled, enter the command with the `-
state```_on_`` parameter.

== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the specified Vserver.

``[-state {on|off}]`` - SnapDiff RPC Server state [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the specified SnapDiff RPC server state.

== Examples

The following example displays information about all the Vservers with
SnapDiff RPC server configured:
```

cluster1::> vserver snapdiff-rpc-server show
Vserver SnapDiff RPC Server State
---------- ------------------------
vs0 on
vs1 off

2 entries were displayed.

```
:leveloffset: -1


= vserver vscan commands

:leveloffset: +1


[[ID5b63546d71624967a20b50c3a390e98]]
= vserver vscan disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable Vscan on a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan disable` command disables Vscan on a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
disable Vscan.

== Examples

The following example disables Vscan on Vserver vs1.
```

cluster1::> vserver vscan disable -vserver vs1

cluster1::> vserver vscan show -vserver vs1
Vserver: vs1
Vscan Status: off

```
[[ID02722c7564d87bb5ed088a07379b7f49]]
= vserver vscan enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable Vscan on a Vserver

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan enable` command enables Vscan on a Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
enable Vscan. The Vscan configuration must already exist.

== Examples

The following example enables Vscan on Vserver vs1.
```

cluster1::> vserver vscan enable -vserver vs1

cluster1::> vserver vscan show -vserver vs1
Vserver: vs1
Vscan Status: on

```
[[IDdab40914e9c52ddd519f63cbe4dfd725]]
= vserver vscan reset
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Discard cached scan information

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan reset` command discards the cached information of the
files that have been successfully scanned. After running this command, the
files are scanned again when they are accessed.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver for which you want to
discard the cached information.

== Examples

The following example discards the cached information of the successfully
scanned files.
```

cluster1::> vserver vscan reset -vserver vs1
Warning: Running this command can cause performance degradation because files are scanned again when they are accessed.
Do you want to continue? {y|n}: y

cluster1::>

```
[[ID54db9f851d044e0721861074f391ff37]]
= vserver vscan show-events
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Vscan events
```

```
*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _advanced_ privilege level.

== Description

The `vserver vscan show-events` command displays contents of the event
log, which is generated by the cluster to capture important events. If you
do not specify any parameters, the command displays the following
information for all Vscan servers:

* Vserver name
* Node name
* Vscan server
* Event type
* Event time



You can specify the `-fields` parameter to specify which fields of
information to display. In addition to the fields above, you can display
the following fields:

* File path
* Vscan server vendor
* Vscan server version
* Disconnect reason
* Scan engine status code
* Vserver LIF used for connection
* Consecutive occurrence count



== 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  ]`` [.advaced-priv]#(privilege: advanced)# }::
If you specify the `-instance` parameter, the command displays detailed
information about all fields.

``[-node {<nodename>|local}]`` - Node [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
```

```
the events that have occurred on the specified node.

``[-vserver <vserver name>]`` - Vserver [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the events that have occurred for the specified Vserver.

``[-event-time <MM/DD/YYYY HH:MM:SS>]`` - Event Log Time [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the events that have occurred at the specified time.

``[-server <IP Address>]`` - Server [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the events that have occurred for the specified server.

``[-event-type <event-type>]`` - Event Type [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the events that are of the specified event type.

``[-file-path <text>]`` - File Path [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the events that have the specified file path.

``[-vendor <text>]`` - Vscanner Vendor [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the events that have the specified scan-engine vendor.

``[-version <text>]`` - Vscanner Version [.advaced-priv]#(privilege:
advanced)#::
If you specify this parameter, the command displays information only about
the events that have the specified scan-engine version.

``[-disconnect-reason <reason>]`` - Server Disconnect Reason [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the events that have the specified reason of the server disconnection.

``[-lif <IP Address>]`` - Vserver LIF Used for Connection [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the events that have the specified IP address, which is used for
```

```
connecting ONTAP with the Vscan server.

== Examples

The following example displays all the events captured in the cluster:
```

cluster1::*> vserver vscan show-events

Vserver Node Server Event Type Event Time
---------- --------------- --------------- ----------------- -----------------
vs1 Cluster-01 192.168.1.1 file-infected 9/5/2014 11:37:38
vs1 Cluster-01 192.168.1.1 scanner-updated 9/5/2014 11:37:08
vs1 Cluster-01 192.168.1.1 scanner-connected 9/5/2014 11:34:55
3 entries were displayed.

```
The following example displays detailed event information about all the
infected files:
```

cluster1::*> vserver vscan show-events -event-type file-infected -instance
Node: Cluster-01
Vserver: vs1
Event Log Time: 9/5/2014 11:37:38
Server: 192.168.1.1
Event Type: file-infected
File Path: \\1
Vscanner Vendor: mighty master anti-evil scanner
Vscanner Version: 1.0
Server Disconnect Reason: -
Vserver LIF Used for Connection: 192.168.41.231

```
[[ID31518a36958d8899deab75600f939c43]]
= vserver vscan show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Vscan status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan show ` command displays Vscan status information of the
Vservers. If you do not specify any parameters, the command displays the
following information about all Vservers:

* Vserver name
* Vscan status



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the specified Vserver.

``[-vscan-status {on|off}]`` - Vscan Status::
If you specify this parameter, the command displays information only about
the Vservers that have the specified status.

== Examples

The following example displays the Vscan status information.
```

cluster1::> vserver vscan show
Vserver Vscan Status
--- ---
vs1 on
vs2 off
2 entries were displayed.

```
[[IDf5bff100ecb23d1202a51d8c12e31503]]
= vserver vscan connection-status show-all
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display Vscan servers connection status

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan connection-status show-all` command displays connection
status information of the external virus-scanning servers, or "Vscan
servers". If you do not specify any parameters, the command displays the
following information for all Vscan servers:

* Vserver name
* Node name
* Vscan server
* Connection status
* Disconnect reason



You can specify the `-fields` parameter to specify which fields of
information to display. In addition to the fields above, you can display
the following fields:

* Server type
* Vscan server vendor
* Vscan server version
* Privileged user
* Vscan server connected since
* Vscan server disconnected since
* Vserver LIF used for connection
```

```
== 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>|local}]`` - Node::
If you specify this parameter, the command displays information only about
the Vscan servers attached to the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the Vscan servers for the specified Vserver.

``[-server <IP Address>]`` - Server::
If you specify this parameter, the command displays information only about
the Vscan server that you specify.

``[-server-status <Status>]`` - Server Status::
If you specify this parameter, the command displays information only about
the Vscan servers that have the specified status.

``[-server-type <Server type>]`` - Server Type::
If you specify this parameter, the command displays information only about
the Vscan servers that have the specified server type.

``[-vendor <text>]`` - Vscanner Vendor::
If you specify this parameter, the command displays information only about
the Vscan servers that are running scan-engine of the specified vendor.

``[-version <text>]`` - Vscanner Version::
If you specify this parameter, the command displays information only about
the Vscan servers that are running scan-engine of the specified version.

``[-disconnect-reason <reason>]`` - Server Disconnect Reason::
If you specify this parameter, the command displays information only about
the Vscan servers that are disconnected because of the specified reason.

``[-disconnected-since <MM/DD/YYYY HH:MM:SS>]`` - Time When Vscanner Was
Disconnected::
```

```
If you specify this parameter, the command displays information only about
the Vscan servers that have been disconnected since the specified time.

``[-privileged-user <text>]`` - Privileged User Used for Connection::
If you specify this parameter, the command displays information only about
the Vscan servers that are connected to ONTAP using the specified
privileged user.

``[-connected-since <MM/DD/YYYY HH:MM:SS>]`` - Time When Vscanner Was
Connected::
If you specify this parameter, the command displays information only about
the Vscan servers that have been connected since the specified time.

``[-lif <IP Address>]`` - Vserver LIF Used for Connection::
If you specify this parameter, the command displays information only about
the Vscan servers that have used the specified IP address for connecting
to ONTAP.

== Examples

The following example displays connection-status information about all
Vscan servers.
```

cluster1::> vserver vscan connection-status show-all
Connection
Vserver Node Server Status Disconnect Reason
---------- ---------- ---------- ---------- ----------
vs1 Cluster-01 1.1.1.1 disconnected remote-closed
vs1 Cluster-01 2.2.2.2 connected -
vs2 Cluster-01 3.3.3.3 disconnected no-data-lif
vs2 Cluster-01 4.4.4.4 disconnected no-data-lif
4 entries were displayed.

```
The following example displays detailed connection-status information
about all Vscan servers which are connected.
```

cluster1::> vserver vscan connection-status show-all -instance
-server-status connected
Node: Cluster-01
Vserver: vs1
Server: 2.2.2.2
Server Status: connected
Server Type: primary
Vscanner Vendor: XYZ
Vscanner Version: 1.12.2
Server Disconnect Reason: -
Time When Server Was Disconnected: -
Privileged User Used for Connection: cifs\u2

Time When Server Was Connected: 6/3/2013 08:44:21
Vserver LIF Used for Connection: 10.238.41.223

```
[[ID742de80b1e40f2ba78dbc301ee1a6a43]]
= vserver vscan connection-status show-connected
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display connection status of connected Vscan servers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan connection-status show-connected` command displays
connection status information of the connected external virus-scanning
servers, or "Vscan servers". If you do not specify any parameters, the
command displays the following information for all Vscan servers:

* Vserver name
* Node name
* Vscan server
* Vscan server vendor
* Privileged user



You can specify the `-fields` parameter to specify which fields of
information to display. In addition to the fields above, you can display
the following fields:

* Server type
* Vscan server version
* Vscan server connected since
* Vserver LIF used for connection



== 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>|local}]`` - Node::
If you specify this parameter, the command displays information only about
the Vscan servers attached to the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the Vscan servers for the specified Vserver.

``[-server <IP Address>]`` - Server::
If you specify this parameter, the command displays information only about
the Vscan server that you specify.

``[-vendor <text>]`` - Vscan Server Vendor::
If you specify this parameter, the command displays information only about
the Vscan servers that are running scan-engine of the specified vendor.

``[-version <text>]`` - Vscan Server Version::
If you specify this parameter, the command displays information only about
the Vscan servers that are running scan-engine of the specified version.

``[-privileged-user <text>]`` - Privileged User Used for Connection::
If you specify this parameter, the command displays information only about
the Vscan servers that are connected to ONTAP using the specified
privileged user.

``[-connected-since <MM/DD/YYYY HH:MM:SS>]`` - Time When Vscan Server Was
Connected::
If you specify this parameter, the command displays information only about
the Vscan servers that have been connected since the specified time.

``[-server-type <Server type>]`` - Server Type::
If you specify this parameter, the command displays information only about
the Vscan servers that have the specified server type.

``[-lif <IP Address>]`` - Vserver LIF Used for Connection::
If you specify this parameter, the command displays information only about
the Vscan servers that have used the specified IP address for connecting
to ONTAP.

== Examples
```

```
The following example displays connection-status information about all
connected Vscan servers.
```

cluster1::> vserver vscan connection-status show-connected
Privileged
Vserver Node Server Vendor User
---------- ---------- ---------- ---------- ----------
vs1 Cluster-01 1.1.1.1 ABC cifs\u2
vs1 Cluster-01 2.2.2.2 XYZ cifs\u2
2 entries were displayed.

```
The following example displays detailed connection-status information
about connected Vscan servers which are running XYZ scan-engine.
```

cluster1::> vserver vscan connection-status show-connected -instance -vendor XYZ
Node: Cluster-01
Vserver: vs1
Server: 2.2.2.2
Vscanner Vendor: XYZ
Vscanner Version: 1.12
Privileged User Used for Connection: cifs\u2
Time When Vscanner Was Connected: 6/3/2013 08:44:21
Server Type: primary
Vserver LIF Used for Connection: 10.238.41.223

```
[[IDc1d26104fcc8204f24263b7b205a1042]]
= vserver vscan connection-status show-not-connected
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display connection status of Vscan servers which are allowed to connect
but not yet connected

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan connection-status show-not-connected` command displays
connection status information of the external virus-scanning servers, or
"Vscan servers" that are ready to accept connection but are not yet
connected. This command could be useful for troubleshooting. If you do not
specify any parameters, the command displays the following information for
```

```
all Vscan servers:

* Vserver name
* Node name
* Vscan server
* Connection status
* Disconnect reason



You can specify the `-fields` parameter to specify which fields of
information to display. In addition to the fields above, you can display
the following fields:

* Server type
* Vscan server disconnected since



== 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>|local}]`` - Node::
If you specify this parameter, the command displays information only about
the Vscan servers attached to the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the Vscan servers for the specified Vserver.

``[-server <IP Address>]`` - Server::
If you specify this parameter, the command displays information only about
the Vscan server that you specify.

``[-server-status <Status>]`` - Server Status::
If you specify this parameter, the command displays information only about
the Vscan servers that have the specified status.

``[-disconnect-reason <reason>]`` - Server Disconnect Reason::
```

```
If you specify this parameter, the command displays information only about
the Vscan servers that are disconnected because of the specified reason.

``[-disconnected-since <MM/DD/YYYY HH:MM:SS>]`` - Time When Vscan Server
Was Disconnected::
If you specify this parameter, the command displays information only about
the Vscan servers that have been disconnected since the specified time.

``[-server-type <Server type>]`` - Server Type::
If you specify this parameter, the command displays information only about
the Vscan servers that have the specified server type.

== Examples

The following example displays connection-status information about all
Vscan servers which are ready to accept connection but not yet connected.
```

cluster1::> vserver vscan connection-status show-not-connected
Connection Disconnect
Vserver Node Server Status Reason
---------- ---------- ---------- ---------- ----------
vs2 Cluster-01 3.3.3.3 disconnected invalid-
session-id
vs2 Cluster-01 4.4.4.4 disconnected remote-closed
2 entries were displayed.

```
The following example displays detailed connection-status information
about Vscan servers which are disconnected because the connection is
remotely closed.
```

cluster1::> vserver vscan connection-status show-not-connected -instance
-disconnect-reason remote-closed
Node: Cluster-01
Vserver: vs2
Server: 4.4.4.4
Server Status: disconnected
Server Disconnect Reason: remote-closed
Time When Vscanner Was Disconnected: 6/4/2013 06:51:32
Server Type: primary

```
[[ID745992a8a8e55ea0fe21325d84224068]]
= vserver vscan connection-status show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display Vscan servers connection status summary

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan connection-status show` command displays connection
status summary of the external virus-scanning servers, or "Vscan servers"
for a Vserver. If you do not specify any parameters, the command displays
the following information for all Vservers:

* Vserver name
* Node name
* List of connected Vscan servers
* Connected count



== 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>|local}]`` - Node::
If you specify this parameter, the command displays information only about
the Vscan servers attached to the specified node.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the Vscan servers for the specified Vserver.

``[-servers <IP Address>,...]`` - List of Connected Vscan Servers::
If you specify this parameter, the command displays information only about
the Vservers that have the specified server or servers.

``[-connection-count <integer>]`` - Number of Connected Vscan Servers
Serving the Vserver::
If you specify this parameter, the command displays information only about
the Vservers that have the specified connection count.
```

```
== Examples

The following example displays connection-status summary for all Vservers.
```

cluster1::> vserver vscan connection-status show
Connected Connected
Vserver Node Server-Count Servers
--- --- --- ---
vs1 Cluster-01 2 1.1.1.1, 2.2.2.2
vs2 Cluster-01 0 -
2 entries were displayed.

```
[[ID5a6595965563c816e185fb808ecbf484]]
= vserver vscan on-access-policy create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an On-Access policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy create` command creates an On-Access
policy.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
create an On-Access policy.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the On-Access policy that you want to
create. An On-Access policy name can be up to 256 characters long and is a
string that can only contain any combination of ASCII-range alphanumeric
characters (a-z, A-Z, 0-9), "_", "-" and ".".

``-protocol <CIFS>`` - File-Access Protocol::
This parameter specifies the protocol name for which the On-Access policy
will be created. Currently only CIFS is supported.
```

```
``[-filters {scan-ro-volume|scan-execute-access}]`` - Filters::
This parameter specifies a list of filters which can be used to define the
scope of the On-Access policy more precisely. The list can include one or
more of the following:
+
  * ``_scan-ro-volume_`` - Enable scans for read-only volume.
  * ``_scan-execute-access_`` - Scan only files opened with execute-access
(CIFS only).



``[-scan-mandatory {on|off}]`` - Mandatory Scan::
This parameter specifies if access to a file is allowed or denied when
there are no external virus-scanning servers available for virus scanning.
The default value for this parameter is "on", this denies file access if
an external virus-scanning server is not available. This parameter has no
impact when an external virus-scanning server is available for file
scanning because access to the file is allowed or denied based on the
response from the virus-scanning server.

``[-max-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Max File Size Allowed
for Scanning::
This parameter specifies the maximum size of the file which will be
considered for virus scanning. By default, it is ``_2GB_`` .

``[-paths-to-exclude <File path>,...]`` - File Paths Not to Scan::
This parameter specifies a list of paths, separated by commas, to exclude
from virus scanning. This path is given from the root of the Vserver and
can be up to 255 characters long. By default, no paths are excluded. CIFS
protocol based On-Access policies must use "\" as the path separator. The
path can be in one of the following forms:
+
  * ``_\dir1\dir2\name_`` - This would match "\dir1\dir2\name" as well as
"\dir1\dir2\name\...".
  * ``_\dir1\dir2\name\_`` - This would only match "\dir1\dir2\name\...".



+
NOTE: If you are using the CLI, you must delimit all paths with double
quotation marks ("). For instance, to add the paths "\vol\a b\" and
"\vol\a,b\" to the `-paths-to-exclude` in the CLI, type ``_"\vol\a
b\","\vol\a,b\"_`` at the command prompt. To add a "?" to the expression,
press ESC followed by the "?".
+
```

```
``[-file-ext-to-exclude <File extension>,...]`` - File Extensions Not to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to exclude from virus scanning. By default, no file extensions are
excluded. Each file extension can be up to 16 characters long. The `-file-
ext-to-exclude` supports wildcard patterns containing "*" and "?". Pattern
matching is defined as:
+
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` would match mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` would
match mp3, mp4 but not mp and mpeg.



+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks (\"). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
+



``[-file-ext-to-include <File extension>,...]`` - File Extensions to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to include for virus scanning. By default it is ``_*_`` , which means all
the file extensions are considered for virus scanning except those which
match one of the patterns provided in `-file-ext-to-exclude` list. Each
file extension can be up to 16 characters long. The `-file-ext-to-include`
supports wildcard patterns containing "*" and "?". Pattern matching is
defined as:
+
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` would match mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` would
match mp3, mp4 but not mp and mpeg.



+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks (\"). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
+

+
```

```
NOTE: If you specify both `-file-ext-to-include` and `-file-ext-to-
exclude` lists, then only those file extensions are considered for virus
scanning which match one of the patterns provided in `-file-ext-to-
include` list but do not match any of the patterns provided in `-file-ext-
to-exclude` list.
+


``[-scan-files-with-no-ext {true|false}]`` - Scan Files with No
Extension::
This parameter specifies if the files without any extension are considered
for virus scanning or not. By default, it is true.

== Examples

The following example creates an On-Access policy.
```

cluster1::> vserver vscan on-access-policy create -vserver vs1 -policy-name test
-protocol CIFS -scan-mandatory on -filters scan-ro-volume -max-file-size 3GB
-file-ext-to-exclude "mp3","txt" -file-ext-to-include "mp*","tx*"
-paths-to-exclude "\vol\a b\","\vol\a,b\"

cluster1::> vserver vscan on-access-policy show -instance -vserver vs1 -policy-name test
Vserver: vs1
Policy: test
Policy Status: off
Policy Config Owner: vserver
File-Access Protocol: CIFS
Filters: scan-ro-volume
Mandatory Scan: on
Max File Size Allowed for Scanning: 3GB
File Paths Not to Scan: \vol\a b\, \vol\a,b\
File Extensions Not to Scan: mp3, txt
File Extensions to Scan: mp*, tx*
Scan Files with No Extension: true

```
[[ID24e91f55208d5410b1485fbc38ef5edc]]
= vserver vscan on-access-policy delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an On-Access policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy delete` command deletes an On-Access
policy.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver from which you want to
delete an On-Access policy.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the On-Access policy that you want to
delete.

== Examples

The following example deletes an On-Access policy.
```

cluster1::> vserver vscan on-access-policy delete -vserver vs1 -policy-name test

cluster1::> vserver vscan on-access-policy show -vserver vs1 -policy-name test
There are no entries matching your query.

```
[[ID92172a5d387c1ff23290b4fad4f53196]]
= vserver vscan on-access-policy disable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Disable an On-Access policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy disable` command disable an On-Access
policy for the specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
disable an On-Access policy. The Vserver administrator can disable On-
Access policies created within the scope of the Vserver and can also
disable an On-Access policy created by the cluster administrator. The
cluster administrator can disable On-Access policies for any Vserver.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the On-Access policy you want to
disable.

== Examples

The following command disable an On-Access policy on specified Vserver.
```

cluster1::> vserver vscan on-access-policy disable -vserver vs1 -policy-name new

cluster1::> vserver vscan on-access-policy show -instance -vserver vs1 -policy-name new
Vserver: vs1
Policy: new
Policy Status: off
Policy Config Owner: vserver
File-Access Protocol: CIFS
Filters: scan-ro-volume
Mandatory Scan: on
Max File Size Allowed for Scanning: 4GB
File-Paths Not to Scan: \vol\temp

File-Extensions Not to Scan: txt

```
[[IDbef71235b1036a4533e7011b5dfe91be]]
= vserver vscan on-access-policy enable
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Enable an On-Access policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy enable` command enables an On-Access
policy for the specified Vserver. Only one On-Access policy of a specific
protocol can be enabled at one time.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
enable an On-Access policy. The Vserver administrator can enable On-Access
policy created within the scope of the Vserver or the cluster. The cluster
administrator can enable On-Access policy for any Vserver but cannot
enable them with a scope of cluster. The scope is determined at a Vserver
level.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the On-Access policy you want to
enable.

== Examples

The following command enables an On-Access policy on specified Vserver.
```

cluster1::> vserver vscan on-access-policy enable -vserver vs1 -policy-name new

cluster1::> vserver vscan on-access-policy show -instance -vserver vs1 -policy-name new
Vserver: vs1
Policy: new
Policy Status: on
Policy Config Owner: vserver
File-Access Protocol: CIFS

Filters: scan-ro-volume
Mandatory Scan: on
Max File Size Allowed for Scanning: 4GB
File-Paths Not to Scan: \vol\temp
File-Extensions Not to Scan: txt

```
[[ID138c9370f650f0b9a6e718cdc2334eef]]
= vserver vscan on-access-policy modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify an On-Access policy

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy modify` command modifies an On-Access
policy.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
modify an On-Access policy.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the On-Access policy that you want to
modify.

``[-filters {scan-ro-volume|scan-execute-access}]`` - Filters::
This parameter specifies a list of filters which can be used to define the
scope of the On-Access policy more precisely. The list can include one or
more of the following:
+
  * ``_scan-ro-volume_`` - Enable scans for read-only volume.
  * ``_scan-execute-access_`` - Scan only files opened with execute-access
(CIFS only).



``[-scan-mandatory {on|off}]`` - Mandatory Scan::
This parameter specifies whether access to a file is allowed if there are
```

```
no external virus-scanning servers available for virus scanning.

``[-max-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Max File Size Allowed
for Scanning::
This parameter specifies the maximum size of the file which will be
considered for virus scanning.

``[-paths-to-exclude <File path>,...]`` - File Paths Not to Scan::
This parameter specifies a list of paths, separated by commas, to exclude
from virus scanning. This path is given from the root of the Vserver and
can be up to 255 characters long. CIFS protocol based On-Access policies
must use "\" as the path separator. The path can be in one of the
following forms:
+
  * ``_\dir1\dir2\name_`` - This would match "\dir1\dir2\name" as well as
"\dir1\dir2\name\...".
  * ``_\dir1\dir2\name\_`` - This would only match "\dir1\dir2\name\...".


+
NOTE: If you are using the CLI, you must delimit all paths with double
quotation marks ("). For instance, to add the paths "\vol\a b\" and
"\vol\a,b\" to the `-paths-to-exclude` in the CLI, type ``_"\vol\a
b\","\vol\a,b\"_`` at the command prompt. To add a "?" to the expression,
press ESC followed by the "?".
+


``[-file-ext-to-exclude <File extension>,...]`` - File Extensions Not to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to exclude from virus scanning. Each file extension can be up to 16
characters long. The `-file-ext-to-exclude` supports wildcard patterns
containing "*" and "?". Pattern matching is defined as:
+
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` would match mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` would
match mp3, mp4 but not mp and mpeg.


+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks ("). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
```

```
+


``[-file-ext-to-include <File extension>,...]`` - File Extensions to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to include for virus scanning. Each file extension can be up to 16
characters long. The `-file-ext-to-include` supports wildcard patterns
containing "*" and "?". Pattern matching is defined as:
+
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` would match mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` would
match mp3, mp4 but not mp and mpeg.


+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks ("). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
+

+
NOTE: If you specify both `-file-ext-to-include` and `-file-ext-to-
exclude` lists, then only those file extensions are considered for virus
scanning which match one of the patterns provided in `-file-ext-to-
include` list but do not match any of the patterns provided in `-file-ext-
to-exclude` list.
+


``[-scan-files-with-no-ext {true|false}]`` - Scan Files with No
Extension::
This parameter specifies if the files without any extension are considered
for virus scanning or not.

== Examples

The following example modifies an On-Access policy.
```

cluster1::> vserver vscan on-access-policy modify -vserver vs1 -policy-name test
-protocol CIFS -scan-mandatory on -filters scan-ro-volume -max-file-size 10GB
-file-ext-to-exclude "mp3" -file-ext-to-include "mp*" -scan-files-with-no-ext false
-paths-to-exclude "\vol1\temp","\vol2\a"

cluster1::> vserver vscan on-access-policy show -instance -vserver vs1 -policy-name test

Vserver: vs1
Policy: test
Policy Status: off
Policy Config Owner: vserver
File-Access Protocol: CIFS
Filters: scan-ro-volume
Mandatory Scan: off
Max File Size Allowed for Scanning: 10GB
File Paths Not to Scan: \vol1\temp, \vol2\a
File Extensions Not to Scan: mp3
File Extensions to Scan: mp*
Scan Files with No Extension: false

```
[[IDfea7b563b6f5d3c0eec10b0e3f6013e5]]
= vserver vscan on-access-policy show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display On-Access policies

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy show` command displays information
about the On-Access policies belonging to the Vserver. It also displays
the current status in Vserver scope. If you do not specify any parameters,
the command displays the following information about all On-Access
policies:

* Vserver name
* Policy name
* Policy status
* Policy owner
* Protocol
* File paths to exclude
* File extensions to exclude



You can specify the `-fields` parameter to specify which fields of
information to display about On-Access policies. In addition to the fields
above, you can display the following fields:
```

```
* List of filters
* Mandatory scan
* Max file size
* File extensions to include
* Scan files without extension



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the On-Access policies for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
the specified On-Access policy.

``[-policy-status {on|off}]`` - Policy Status::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified status.

``[-owner <Configuration owner>]`` - Policy Config Owner::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified owner.

``[-protocol <CIFS>]`` - File-Access Protocol::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified protocol.

``[-filters {scan-ro-volume|scan-execute-access}]`` - Filters::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified filter or filters in the
filter list.

``[-scan-mandatory {on|off}]`` - Mandatory Scan::
If you specify this parameter, the command displays information only about
the On-Access policies that have mandatory scanning enabled.
```

```
``[-max-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Max File Size Allowed
for Scanning::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified max-file-size.

``[-paths-to-exclude <File path>,...]`` - File Paths Not to Scan::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified path or paths in the
exclude list.

``[-file-ext-to-exclude <File extension>,...]`` - File Extensions Not to
Scan::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified file extension or
extensions in the exclude list.

``[-file-ext-to-include <File extension>,...]`` - File Extensions to
Scan::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified file extension or
extensions in the include list.

``[-scan-files-with-no-ext {true|false}]`` - Scan Files with No
Extension::
If you specify this parameter, the command displays information only about
the On-Access policies that have the specified value.

== Examples

The following example displays information about all On-Access policies.
```

cluster1::> vserver vscan on-access-policy show
Policy Policy File-Ext Policy
Vserver Name Owner Protocol Paths Excluded Excluded Status
---------- --------- -------- -------- ---------- ---------- ------
Cluster default_ cluster CIFS - - off
CIFS
vs1 default_ cluster CIFS - - on
CIFS
vs1 new vserver CIFS \vol\temp txt off
vs2 default_ cluster CIFS - - on
CIFS
4 entries were displayed.

```
The following example displays detailed information about an On-Access
policy.
```

cluster1::> vserver vscan on-access-policy show -instance -vserver vs1 -policy-name new
Vserver: vs1
Policy: new
Policy Status: off
Policy Config Owner: vserver
File-Access Protocol: CIFS
Filters: scan-ro-volume
Mandatory Scan: on
Max File Size Allowed for Scanning: 4GB
File Paths Not to Scan: \vol\temp
File Extensions Not to Scan: txt
File Extensions to Scan: *
Scan Files with No Extension: true

```
[[ID921dadc220cb3ada2501f3871b64ed71]]
= vserver vscan on-access-policy file-ext-to-exclude add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to the list of file extensions to exclude

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy file-ext-to-exclude add` command adds
a file extension or a list of file extensions that must be excluded from
scanning to the specified policy name.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
on-access policy to which you want to add a file extension or a list of
file extensions that must be excluded from scanning.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the on-access policy to which you
want to add a file extension or a list of file extensions that must be
excluded from scanning.

``-file-ext-to-exclude <File extension>,...`` - File Extensions Not to
Scan::
This parameter specifies the file extension or a list of file extensions
that must be excluded from scanning.

== Examples

The following example adds a list of file extensions that must be excluded
from scanning to the specified on-access policy:
```

cluster1::> vserver vscan on-access-policy file-ext-to-exclude add -vserver vs1
-policy-name policy1 -file-ext-to-exclude txt,mp4

cluster1::> vserver vscan on-access-policy file-ext-to-exclude show -vserver vs1
-policy-name policy1

Vserver: vs1
Policy: policy1
File-Extensions Not to Scan: mp3, mp4, txt, wav

```
[[ID659f593af0367bde29452f88fd45d6b7]]
= vserver vscan on-access-policy file-ext-to-exclude remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from the list of file extensions to exclude

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy file-ext-to-exclude remove` command
removes a file extension or a list of file extensions that are excluded
from scanning from the specified policy name.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
on-access policy from which you want to remove a file extension or a list
of file extensions that are excluded from scanning.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the on-access policy from which you
want to remove a file extension or a list of file extensions that are
excluded from scanning.

``-file-ext-to-exclude <File extension>,...`` - File Extensions Not to
Scan::
This parameter specifies the file extension or a list of file extensions
that must be removed from the on-access policy.

== Examples

The following example removes a list of file extensions that are to be
excluded from scanning from the specified on-access policy:
```

cluster1::> vserver vscan on-access-policy file-ext-to-exclude remove -vserver vs1

-policy-name policy1 -file-ext-to-exclude mp3,txt

cluster1::> vserver vscan on-access-policy file-ext-to-exclude show -vserver vs1
-policy-name policy1
Vserver: vs1
Policy: policy1
File-Extensions Not to Scan: mp4, wav

```
[[IDd8cbe5f705f82a05197ba0a67ee0fafd]]
= vserver vscan on-access-policy file-ext-to-exclude show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of file extensions to exclude

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy file-ext-to-exclude show` command
displays the list of file extensions that are excluded from scanning
belonging to the Vserver. If you do not specify any parameters, the
command displays the following information about all on-access policies:

* Vserver name
* Policy name
* List of File-Extensions to exclude



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
```

```
the policy names for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
the specified policy name.

``[-file-ext-to-exclude <File extension>,...]`` - File Extensions Not to
Scan::
If you specify this parameter, the command displays information only about
the policies that have the specified file extensions that are excluded
from scanning.

== Examples

The following example displays the list of file extensions that are
excluded from scanning for all the policies:
```

cluster1::> vserver vscan on-access-policy file-ext-to-exclude show
Vserver Policy Name File-Ext Excluded
--- --- ---
cluster1 default_CIFS txt
vs1 default_CIFS txt
vs1 policy1 mp4, wav
vs1 policy3 wmv
vs2 default_CIFS txt
vs2 policy2 mp3
6 entries were displayed.

```
[[ID5d245e70188d2bfa38919f636401ad20]]
= vserver vscan on-access-policy file-ext-to-include add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to the list of file extensions to include

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy file-ext-to-include add` command adds
a file extension or list of file extensions to include for virus scanning
to the specified policy.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
on-access policy to which you want to add a file extension or a list of
file extensions to include for virus scanning.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the on-access policy to which you
want to add a file extension or a list of file extensions to include for
virus scanning.

``-file-ext-to-include <File extension>,...`` - File Extensions to Scan::
This parameter specifies the file extension or a list of file extensions
to include for virus scanning.

== Examples

The following example adds a list of file extensions to include for virus
scanning to the specified on-access policy.
```

cluster1::> vserver vscan on-access-policy file-ext-to-include add -vserver vs1
-policy-name policy1 -file-ext-to-include "mp*","tx*"

cluster1::> vserver vscan on-access-policy file-ext-to-include show -vserver vs1
-policy-name policy1
Vserver: vs1

Policy: policy1
File Extensions to Scan: mp*, tx*, wav

```
[[IDe46897013e706d5538dd501917bcc44c]]
= vserver vscan on-access-policy file-ext-to-include remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from the list of file extensions to include

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy file-ext-to-include remove` command
removes a file extension or list of file extension that are included for
virus scanning from the specified policy.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
on-access policy from which you want to remove a file extension or list of
file extensions that are included for virus scanning.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the on-access policy from which you
want to remove a file extension or a list of file extensions that are
included for virus scanning.

``-file-ext-to-include <File extension>,...`` - File Extensions to Scan::
This parameter specifies the file extension or a list of file extensions
that you want to remove from the specified on-access policy.

== Examples

The following example removes a list of file extensions from the specified
on-access policy.
```

cluster1::> vserver vscan on-access-policy file-ext-to-include remove -vserver vs1
-policy-name policy1 -file-ext-to-include "txt*,"wav"

cluster1::> vserver vscan on-access-policy file-ext-to-include show -vserver vs1
-policy-name policy1
Vserver: vs1
Policy: policy1
File Extensions to Scan: mp*

```
[[ID2a64208cdb645737ece70d0926aaf8fd]]
= vserver vscan on-access-policy file-ext-to-include show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of file extensions to include

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy file-ext-to-include show` command
displays the list of file extensions to include for virus scanning
belonging to the Vserver. If you do not specify any parameters, the
command displays the following information about all on access policies:

* Vserver name
* Policy name
* List of File-Extensions to Scan



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the policies for the specified Vserver.
```

```
``[-policy-name <Policy name>]`` - Policy::
If you specify this parameter, the command displays information only about
the specified policy.

``[-file-ext-to-include <File extension>,...]`` - File Extensions to
Scan::
If you specify this parameter, the command displays information only about
the policies that have the specified file extensions that are included for
virus scanning.

== Examples

The following example displays the list of file extensions that are
included for virus scanning for all policies.
```

cluster1::> vserver vscan on-access-policy file-ext-to-include show
Vserver Policy Name File-Ext Included
--- --- ---
cluster1 default_CIFS *
vs1 default_CIFS *
vs1 policy1 mp*
vs1 policy3 doc*, xl*
vs2 default_CIFS *
vs2 policy2 d*, m*, t*
6 entries were displayed.

```
[[ID77771f329c3d7064063e6364389e3dda]]
= vserver vscan on-access-policy paths-to-exclude add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to the list of paths to exclude

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy paths-to-exclude add` command adds a
path or a list of paths that must be excluded from scanning to the
specified policy name.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
on-access policy to which you want to add a path or a list of paths that
must be excluded from scanning.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the on-access policy to which you
want to add a path or a list of paths that must be excluded from scanning.

``-paths-to-exclude <File path>,...`` - Paths Not to Scan::
This parameter specifies the path or list of paths that must be excluded
from scanning.

== Examples

The following example adds a list of paths that must be excluded from
scanning to the specified on-access policy:
```

cluster1::> vserver vscan on-access-policy paths-to-exclude add -vserver vs1
-policy-name policy1 -paths-to-exclude \test\test2,\test\test3

cluster1::> vserver vscan on-access-policy paths-to-exclude show -vserver vs1
-policy-name policy1
Vserver: vs1
Policy: policy1

File-Paths Not to Scan: \test\test1, \test\test2, \test\test3

```
[[ID3e545d4d9bf53258f1582e33d8efda07]]
= vserver vscan on-access-policy paths-to-exclude remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from the list of paths to exclude

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy paths-to-exclude remove` command
removes a path or a list of paths that are excluded from scanning from the
specified policy name.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
on-access policy from which you want to remove a path or list of paths
that are excluded from scanning.

``-policy-name <Policy name>`` - Policy::
This parameter specifies the name of the on-access policy from which you
want to remove a path or a list of paths that are excluded from scanning.

``-paths-to-exclude <File path>,...`` - Paths Not to Scan::
This parameter specifies the path or a list of paths that must be removed
from the on-access policy.

== Examples

The following example removes a list of paths that are excluded from
scanning from the specified policy name:
```

cluster:> vserver vscan on-access-policy paths-to-exclude remove -vserver vs1
-policy-name policy1 -paths-to-exclude \test\test2,\test\test3

cluster1:> vserver vscan on-access-policy paths-to-exclude show -vserver vs1
-policy-name policy1

Vserver: vs1
Policy: policy1
File-Paths Not to Scan: \test\test1

```
[[IDab921f5e73bc6f4206c17d00178cc5c0]]
= vserver vscan on-access-policy paths-to-exclude show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of paths to exclude

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-access-policy paths-to-exclude show` command
displays the list of paths that are excluded from scanning belonging to
the Vserver. If you do not specify any parameters, the command displays
the following information about all on-access policies:

* Vserver name
* Policy name
* List of Paths to exclude



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the policy names for the specified Vserver.

``[-policy-name <Policy name>]`` - Policy::
```

```
If you specify this parameter, the command displays information only about
the specified policy name.

``[-paths-to-exclude <File path>,...]`` - File Paths Not to Scan::
If you specify this parameter, the command displays information only about
the policies that have the specified paths that are excluded from
scanning.

== Examples

The following example displays the list of paths that are excluded from
scanning for all the policies:
```

```
cluster1::> vserver vscan on-access-policy paths-to-exclude show
Vserver Policy Name Paths Excluded
--------------- ---------------- ------------------------------------------
cluster1 default_CIFS \test\test1
vs1 default_CIFS \test\test1
vs1 policy1 \test\test2,\test\test3
vs1 policy3 \test\test4
vs2 default_CIFS \test\test1
vs2 policy2 \test\test5
6 entries were displayed.
```

```
[[ID505c981163d8a658c6f35cbe13315540]]
= vserver vscan on-demand-task create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create an On-Demand task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task create` command creates an On-Demand
task. The On-Demand task consists of a set of attributes that are used for
configuring the scope of scanning. It also specifies the cron schedule at
which the task should run.

== Parameters

``-vserver <vserver name>`` - Vserver::
```

```
This parameter specifies the name of the Vserver on which you want to
create an On-Demand task.

``-task-name <text>`` - Task Name::
This parameter specifies the name of the On-Demand task that you want to
create. An On-Demand task name can be up to 256 characters long.

``-scan-paths <text>,...`` - List of Scan Paths::
This parameter specifies a list of paths, separated by commas, for virus
scanning. This path is given from the root of the Vserver using UNIX path
delimiter "/".

``-report-directory <text>`` - Report Directory Path::
This parameter specifies a directory path where the On-Demand report file
is created. Each run for a task creates a new file. The report directory
path is given from the root of the Vserver using UNIX path delimiter "/".

``[-schedule <text>]`` - Job Schedule::
This parameter specifies the already existing cron schedule. The On-Demand
task triggers virus scanning for the specified scan-paths at the time
configured in the schedule.
+
NOTE: A Vserver can have only one scheduled task at a time.
+



``[-max-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Max File Size Allowed
for Scanning::
This parameter specifies the maximum size of the file that will be
considered for virus scanning. By default, it is ``_10GB_`` .

``[-paths-to-exclude <text>,...]`` - File Paths Not to Scan::
This parameter specifies a list of paths, separated by commas, to exclude
from virus scanning. This path is given from the root of the Vserver using
UNIX path delimiter "/". By default, no paths are excluded. The path can
be in one of the following forms:
+
  * ``_/dir1/dir2/name_`` - This would match "/dir1/dir2/name" as well as
"/dir1/dir2/name/...".
  * ``_/dir1/dir2/name/_`` - This would only match "/dir1/dir2/name/...".



+
NOTE: If you are using the CLI, you must delimit all paths with double
quotation marks ("). For instance, to add the paths "/vol/a b/" and
"/vol/a,b/" to the `-paths-to-exclude` in the CLI, type ``_"/vol/a
b/","/vol/a,b/"_`` at the command prompt.
```

```
+


``[-file-ext-to-exclude <File extension>,...]`` - File Extensions Not to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to exclude from virus scanning. By default, no file extensions are
excluded. Each file extension can be up to 16 characters long. The `-file-
ext-to-exclude` supports wildcard patterns containing "*" and "?". Pattern
matching is defined as:
+
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` matches mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` matches
mp3, mp4 but not mp and mpeg.


+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks ("). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
+


``[-file-ext-to-include <File extension>,...]`` - File Extensions to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to include for virus scanning. By default it is ``_*_`` , which means all
the file extensions are considered for virus scanning except those that
match one of the patterns provided in `-file-ext-to-exclude` list. Each
file extension can be up to 16 characters long. The `-file-ext-to-include`
supports wildcard patterns containing "*" and "?". Pattern matching is
defined as:
+
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` matches mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` matches
mp3, mp4 but not mp and mpeg.


+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks ("). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
```

```
+

+
NOTE: If you specify both `-file-ext-to-include` and `-file-ext-to-
exclude` lists, then only those file extensions are considered for virus
scanning which match one of the patterns provided in `-file-ext-to-
include` list but do not match any of the patterns provided in `-file-ext-
to-exclude` list.
+


``[-scan-files-with-no-ext {true|false}]`` - Scan Files with No
Extension::
This parameter specifies if the files without any extension are considered
for virus scanning or not. By default, it is true.

``[-request-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Request Service Timeout::
This parameter specifies the timeout value for a scan request. It is used
to specify the time interval in which the node waits for a response from
the Vscan server. Beyond this timeout period, the scan request is
considered as failed. The value for this field must be between 10s and 1h.
By default, it is 5m.

``[-cross-junction {true|false}]`` - Cross Junction::
This parameter specifies if the On-Demand task is allowed to cross volume
junctions. If the parameter is set to false, crossing junctions is not
allowed. By default, it is true.

``[-directory-recursion {true|false}]`` - Directory Recursion::
This parameter specifies if the On-Demand task is allowed to recursively
scan through sub-directories. If the parameter is set to false, recursive
scanning is not allowed. By default, it is true.

``[-scan-priority {low|normal}]`` - Scan Priority::
This parameter specifies the priority of the On-Demand scan requests
generated by this task compared to On-Access scan requests. By default, it
is low.

``[-report-log-level {verbose|info|error}]`` - Report Log Level::
This parameter specifies the log level of the On-Demand report. By
default, it is info.

``[-report-expiry-time
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Expiration Time
for Report::
```

```
This parameter specifies the expiration time for the reports generated by
On-Demand scans. Once this time elapses, the reports are auto-deleted. The
default value is 0, which means reports are retained until they are
manually deleted.

== Examples

The following example creates an On-Demand task:
```

cluster1::> vserver vscan on-demand-task create -vserver vs1 -task-name t1
-scan-paths "/vol1/","/vol2/cifs/" -report-directory "/report"
-schedule daily -max-file-size 5GB -paths-to-exclude "/vol1/cold-files/"
-file-ext-to-include "vmdk?","mp*" -file-ext-to-exclude "mp3","mp4"
-scan-files-with-no-ext false -request-timeout 2m -cross-junction false
-directory-recursion true -scan-priority low -report-log-level verbose
-report-expiry-time 12h
[Job 126]: Vscan On-Demand job is queued. Use the "job show -id 126" command to view the status.

cluster1::> vserver vscan on-demand-task show -instance -vserver vs1 -task-name t1
Vserver: vs1
Task Name: t1
List of Scan Paths: /vol1/, /vol2/cifs/
Report Directory Path: /report
Job Schedule: daily
Max File Size Allowed for Scanning: 5GB
File Paths Not to Scan: /vol1/cold-files/
File Extensions Not to Scan: mp3, mp4
File Extensions to Scan: vmdk, mp*
Scan Files with No Extension: false
Request Service Timeout: 2m
Cross Junction: false
Directory Recursion: true
Scan Priority: low
Report Log Level: verbose

```
[[ID362ccb21db6e253b0906fe2856a85ee9]]
= vserver vscan on-demand-task delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an On-Demand task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task delete` command deletes an On-Demand
task.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver from which you want to
delete an On-Demand task.

``-task-name <text>`` - Task Name::
This parameter specifies the name of the On-Demand task that you want to
delete.

== Examples

The following example deletes an On-Demand task:
```

cluster1::> vserver vscan on-demand-task delete -vserver vs1 -task-name t1

cluster1::> vserver vscan on-demand-task show -vserver vs1 -task-name t1
There are no entries matching your query.

```
[[ID5b2cab58643110feeed53b9de6963696]]
= vserver vscan on-demand-task modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
```

```
Modify an On-Demand task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task modify` command modifies an On-Demand
task. The On-Demand task consists of a set of attributes that are used for
configuring the scope of scanning. It also specifies the cron schedule at
which the task should run.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
modify an On-Demand task.

``-task-name <text>`` - Task Name::
This parameter specifies the name of the On-Demand task that you want to
modify.

``[-scan-paths <text>,...]`` - List of Scan Paths::
This parameter specifies a list of paths, separated by commas, for virus
scanning. This path is given from the root of the Vserver using UNIX path
delimiter "/".

``[-report-directory <text>]`` - Report Directory Path::
This parameter specifies a directory path where the On-Demand report file
is created. Each run for a task creates a new file. The report directory
path is given from the root of the Vserver using UNIX path delimiter "/".

``[-schedule <text>]`` - Job Schedule::
This parameter specifies the already existing cron schedule. The On-Demand
task triggers virus scanning for the specified scan-paths at the time
configured in the schedule. Providing empty schedule ("") unschedules the
task.
+
NOTE: A Vserver can have only one scheduled task at a time.
+


``[-max-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Max File Size Allowed
for Scanning::
This parameter specifies the maximum size of the file which will be
considered for virus scanning.
```

```
``[-paths-to-exclude <text>,...]`` - File Paths Not to Scan::
This parameter specifies a list of paths, separated by commas, to exclude
from virus scanning. This path is given from the root of the Vserver using
UNIX path delimiter "/". The path can be in one of the following forms:
+
  * ``_/dir1/dir2/name_`` - This would match "/dir1/dir2/name" as well as
"/dir1/dir2/name/...".
  * ``_/dir1/dir2/name/_`` - This would only match "/dir1/dir2/name/...".


+
NOTE: If you are using the CLI, you must delimit all paths with double
quotation marks ("). For instance, to add the paths "/vol/a b/" and
"/vol/a,b/" to the `-paths-to-exclude` in the CLI, type ``_"/vol/a
b/","/vol/a,b/"_`` at the command prompt.
+


``[-file-ext-to-exclude <File extension>,...]`` - File Extensions Not to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to exclude from virus scanning. Each file extension can be up to 16
characters long. The `-file-ext-to-exclude` supports wildcard patterns
containing "*" and "?". Pattern matching is defined as:
+
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` matches mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` matches
mp3, mp4 but not mp and mpeg.


+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks ("). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
+


``[-file-ext-to-include <File extension>,...]`` - File Extensions to
Scan::
This parameter specifies a list of file extensions, separated by commas,
to include for virus scanning. Each file extension can be up to 16
characters long. The `-file-ext-to-include` supports wildcard patterns
containing "*" and "?". Pattern matching is defined as:
+
```

```
  * ``_*_`` - Matches any string, including the empty string. For example,
``_mp*_`` matches mp, mp3, mp4, mpeg etc.
  * ``_?_`` - Matches any single character. For example, ``_mp?_`` matches
mp3, mp4 but not mp and mpeg.



+
NOTE: If you are using the CLI, you must delimit all patterns with double
quotation marks ("). For instance, to enter the pattern ``_mp*_`` in the
CLI, type ``_"mp*"_`` at the command prompt. To add a "?" to the
expression, press ESC followed by the "?".
+

+
NOTE: If you specify both `-file-ext-to-include` and `-file-ext-to-
exclude` lists, then only those file extensions are considered for virus
scanning which match one of the patterns provided in `-file-ext-to-
include` list but do not match any of the patterns provided in `-file-ext-
to-exclude` list.
+



``[-scan-files-with-no-ext {true|false}]`` - Scan Files with No
Extension::
This parameter specifies if the files without any extension are considered
for virus scanning or not.

``[-request-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Request Service Timeout::
This parameter specifies the timeout value for a scan request. It is used
to specify the time interval in which the node waits for a response from
the Vscan server. Beyond this timeout period, the scan request is
considered as failed. The value for this field must be between 10s and 1h.

``[-cross-junction {true|false}]`` - Cross Junction::
This parameter specifies if the On-Demand task is allowed to cross volume
junctions. If the parameter is set to false, crossing junctions is not
allowed.

``[-directory-recursion {true|false}]`` - Directory Recursion::
This parameter specifies if the On-Demand task is allowed to recursively
scan through sub-directories. If the parameter is set to false, recursive
scanning is not allowed.

``[-scan-priority {low|normal}]`` - Scan Priority::
This parameter specifies the priority of the On-Demand scan requests
generated by this task compared to On-Access scan requests.
```

```
``[-report-log-level {verbose|info|error}]`` - Report Log Level::
This parameter specifies the log level of the On-Demand report.

``[-report-expiry-time
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Expiration Time
for Report::
This parameter specifies the expiration time for the reports generated by
On-Demand scans. Once this time elapses, the reports are auto-deleted. The
default value is 0, which means reports are retained until they are
manually deleted.

== Examples

The following example modifies an On-Demand task:
```

cluster1::> vserver vscan on-demand-task modify -vserver vs1 -task-name t1
-scan-paths "/vol3/","/vol4/cifs/" -report-directory "/report-dir"
-schedule custom -max-file-size 2GB -paths-to-exclude "/vol1/cold-files/"
-file-ext-to-include "*" -file-ext-to-exclude "mp3","mp4"
-scan-files-with-no-ext true -request-timeout 1m -cross-junction true
[Job 136]: Vscan On-Demand job is queued. Use the "job show -id 136" command to view the status.

cluster1::> vserver vscan on-demand-task show -instance -vserver vs1 -task-name t1
Vserver: vs1
Task Name: t1
List of Scan Paths: /vol3/, /vol4/cifs/
Report Directory Path: /report-dir
Job Schedule: custom
Max File Size Allowed for Scanning: 2GB
File Paths Not to Scan: /vol1/cold-files/
File Extensions Not to Scan: mp3, mp4
File Extensions to Scan: *
Scan Files with No Extension: true
Request Service Timeout: 1m
Cross Junction: true
Directory Recursion: true
Scan Priority: low
Report Log Level: verbose

```
[[ID8e66d134865adc7f2eae1c3a1d89dacc]]
= vserver vscan on-demand-task run
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Run an On-Demand task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task run` command start virus scanning
immediately for an On-Demand task.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
start start virus scanning.

``-task-name <text>`` - Task Name::
This parameter specifies the name of the On-Demand task that you want to
start virus scanning.

== Examples

The following example starts virus scanning an On-Demand task:
```

cluster1::> vserver vscan on-demand-task run -vserver vs1 -task-name t1
[Job 161]: Vscan On-Demand job is queued. Use the "job show -id 161" command to view the status.

```
[[ID61a0e1ea2ea1b1716cd7153d6035d935]]
= vserver vscan on-demand-task schedule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Schedule an On-Demand task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task schedule` command schedules an On-Demand
task.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
schedule an On-Demand task.

``-task-name <text>`` - Task Name::
This parameter specifies the name of the On-Demand task that you want to
schedule.

``-schedule <text>`` - Schedule Name::
This parameter specifies the already existing cron schedule. The On-Demand
task triggers virus scanning for the specified scan-paths at the time
configured in the schedule.
+
NOTE: A Vserver can have only one scheduled task at a time.
+


== Examples

The following example schedules an On-Demand task:
```

cluster1::> vserver vscan on-demand-task schedule -vserver vs1 -task-name t1 -schedule daily
[Job 150]: Vscan On-Demand job is queued. Use the "job show -id 150" command to view the status.

cluster1::> vserver vscan on-demand-task show -instance -vserver vs1 -task-name t1
Vserver: vs1

Task Name: t1
List of Scan Paths: /test
Report Directory Path: /report
Job Schedule: daily
Max File Size Allowed for Scanning: 2GB
File Paths Not to Scan: /vol1/cold-files/
File Extensions Not to Scan: mp3, mp4
File Extensions to Scan: *
Scan Files with No Extension: true
Request Service Timeout: 1m
Cross Junction: true
Directory Recursion: true
Scan Priority: low
Report Log Level: verbose

```
[[IDed7011cdaf6b9c462617883c6324e6bd]]
= vserver vscan on-demand-task show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display On-Demand tasks

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task show` command displays information about
the On-Demand tasks belonging to the Vserver. If you do not specify any
parameters, the command displays the following information about all On-
Demand tasks:

* Vserver name
* Task name
* Scan paths
* Report directory path
* Schedule



You can specify the `-fields` parameter to specify which fields of
information to display about On-Demand tasks. In addition to the fields
above, you can display the following fields:

* Max file size
```

```
* File paths to exclude
* File extensions to exclude
* File extensions to include
* Scan files without extension
* Scan timeout
* Cross junction
* Directory recursion
* Scan priority
* Report log level



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the On-Demand tasks for the specified Vserver.

``[-task-name <text>]`` - Task Name::
If you specify this parameter, the command displays information only about
the specified On-Demand task.

``[-scan-paths <text>,...]`` - List of Scan Paths::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified path or paths in the scan-
paths list.

``[-report-directory <text>]`` - Report Directory Path::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified report-directory.

``[-schedule <text>]`` - Job Schedule::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified schedule.

``[-max-file-size {<integer>[KB|MB|GB|TB|PB]}]`` - Max File Size Allowed
for Scanning::
If you specify this parameter, the command displays information only about
```

```
the On-Demand tasks that have the specified max-file-size.

``[-paths-to-exclude <text>,...]`` - File Paths Not to Scan::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified path or paths in the exclude
list.

``[-file-ext-to-exclude <File extension>,...]`` - File Extensions Not to
Scan::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified file extension or extensions
in the exclude list.

``[-file-ext-to-include <File extension>,...]`` - File Extensions to
Scan::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified file extension or extensions
in the include list.

``[-scan-files-with-no-ext {true|false}]`` - Scan Files with No
Extension::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified value.

``[-request-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Request Service Timeout::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified request-timeout.

``[-cross-junction {true|false}]`` - Cross Junction::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified value.

``[-directory-recursion {true|false}]`` - Directory Recursion::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified value.

``[-scan-priority {low|normal}]`` - Scan Priority::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified scan-priority.

``[-report-log-level {verbose|info|error}]`` - Report Log Level::
If you specify this parameter, the command displays information only about
the On-Demand tasks that have the specified report-log-level.

``[-report-expiry-time
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Expiration Time
```

```
for Report::
This parameter specifies the expiration time for the reports generated by
On-Demand scans. Once this time elapses, the reports are auto-deleted. The
default value is 0, which means reports are retained until they are
manually deleted.

== Examples

The following example displays information about all On-Demand tasks:
```

cluster1::> vserver vscan on-demand-task show
Report
Vserver Task Name Scan Paths Directory Path Schedule
---------- --------- ---------- ---------- ----------
vs1 t1 /test /report -
vs2 t2 /, /test/ /report daily
2 entries were displayed.

```
The following example displays detailed information about an On-Demand
task:
```

cluster1::> vserver vscan on-demand-task show -instance -vserver vs1 -task-name t1
Vserver: vs1
Task Name: t1
List of Scan Paths: /test
Report Directory Path: /report
Job Schedule: -
Max File Size Allowed for Scanning: 2GB
File Paths Not to Scan: /vol1/cold-files/
File Extensions Not to Scan: mp3, mp4
File Extensions to Scan: *
Scan Files with No Extension: true
Request Service Timeout: 1m
Cross Junction: true
Directory Recursion: true
Scan Priority: low
Report Log Level: verbose

```
[[ID06efaa6584cb40b5a2f47de77ab00077]]
= vserver vscan on-demand-task unschedule
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Unschedule an On-Demand task

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task unschedule` command unschedules an On-
Demand task.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
unschedule an On-Demand task.

``-task-name <text>`` - Task Name::
This parameter specifies the name of the On-Demand task that you want to
unschedule.

== Examples

The following example unschedules an On-Demand task:
```

cluster1::> vserver vscan on-demand-task unschedule -vserver vs1 -task-name t1

cluster1::> vserver vscan on-demand-task show -instance -vserver vs1 -task-name t1
Vserver: vs1
Task Name: t1
List of Scan Paths: /test
Report Directory Path: /report
Job Schedule: -
Max File Size Allowed for Scanning: 2GB
File Paths Not to Scan: /vol1/cold-files/
File Extensions Not to Scan: mp3, mp4
File Extensions to Scan: *
Scan Files with No Extension: true
Request Service Timeout: 1m
Cross Junction: true

Directory Recursion: true
Scan Priority: low
Report Log Level: verbose

```
[[ID9f4779ad9be0045ec6f2107348ac5d3f]]
= vserver vscan on-demand-task report delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete an On-Demand report

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task report delete` command deletes an On-
Demand report.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver from which you want to
delete an On-Demand report.

``-task-name <text>`` - Task Name::
This parameter specifies the name of the On-Demand task whose report you
want to delete.

``-report-file <text>`` - Report File Path::
This parameter specifies the path of the report-file whose report record
you want to delete.

``[-delete-report-file {true|false}]`` - Delete Report File Also::
This parameter specifies if the corresponding report file is also to be
deleted. By default, it is false.

== Examples

The following example deletes only On-Demand report record:
```

cluster1::> vserver vscan on-demand-task report delete -vserver vs1 -task-name t1
-report-file /rep/avod_146_20150902_161439.log

cluster1::> vserver vscan on-demand-task report delete -vserver vs1 -task-name t1 -report-file /rep/avod_146_20150902_161439.log
There are no entries matching your query.

```
The following example deletes an On-Demand report file along with the
report record:
```

cluster1::> vserver vscan on-demand-task report delete -vserver vs1 -task-name t1 -report-file /rep/avod_146_20150902_161439.log -delete-report-file true

cluster1::> vserver vscan on-demand-task report delete -vserver vs1 -task-name t1 -report-file /rep/avod_146_20150902_161439.log -delete-report-file true
There are no entries matching your query.

```
[[IDcf03197cd7488e7e6abb18585b67fcf3]]
= vserver vscan on-demand-task report show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display On-Demand reports

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan on-demand-task report show` command displays
information about the On-Demand reports belonging to the Vserver. A new
report record is generated at the end of an On-Demand task run. If you do
not specify any parameters, the command displays the following information
about all On-Demand tasks:

* Vserver name
* Task name
* Report file path
* Number of clean files
* Number of infected files



You can specify the `-fields` parameter to specify which fields of
information to display about On-Demand report. In addition to the fields
above, you can display the following fields:
```

```
* Job ID
* Job duration
* Number of attempted scans
* Number of files skipped from scanning
* Number of already scanned files
* Number of successful scans
* Number of failed scans
* Number of timed-out scans
* Job start time
* Job end time



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the On-Demand reports for the specified Vserver.

``[-task-name <text>]`` - Task Name::
If you specify this parameter, the command displays information only about
the On-Demand reports for the specified task.

``[-report-file <text>]`` - Report File Path::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified report file-path.

``[-job-id <integer>]`` - Job ID::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified job ID.

``[-job-duration <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` -
Job Duration::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-attempted-scans <integer>]`` - Number of Attempted Scans::
```

```
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-skipped-scans <integer>]`` - Number of Files Skipped from Scanning::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-already-scanned-files <integer>]`` - Number of Already Scanned Files::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-successful-scans <integer>]`` - Number of Successful Scans::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-failed-scans <integer>]`` - Number of Failed Scans::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-timedout-scans <integer>]`` - Number of Timedout Scans::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-files-cleaned <integer>]`` - Number of Clean Files::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-files-infected <integer>]`` - Number of Infected Files::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-internal-error <integer>]`` - Number of Internal Error [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-scan-retries <integer>]`` - Number of Scan Retries [.advaced-
priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-job-start-time <MM/DD/YYYY HH:MM:SS>]`` - Job Start Time::
If you specify this parameter, the command displays information only about
the On-Demand reports that have the specified value.

``[-job-end-time <MM/DD/YYYY HH:MM:SS>]`` - Job End Time::
If you specify this parameter, the command displays information only about
```

```
the On-Demand reports that have the specified value.

== Examples

The following example displays information about all On-Demand reports:
```

cluster1::> vscan on-demand-task report show
Files Files
Vserver Task Name Report File Path Cleaned Infected
---------- ---------- ---------- --------- ---------
vs1 t1 /rep/avod_146_20150902_161439.log 6240 5
vs1 t1 /rep/avod_149_20150903_160313.log 115 0
2 entries were displayed.

```
The following example displays detailed information about an On-Demand
task:
```

cluster1::> vscan on-demand-task report show -vserver vs1 -task-name t1
-report-file /rep/avod_146_20150902_161439.log
Vserver: vs1
Task Name: t1
Report File Path: /rep/avod_146_20150902_161439.log
Job ID: 146
Job Duration: 76s
Number of Attempted Scans: 6245
Number of Files Skipped from Scanning: 1286
Number of Already Scanned Files: 987
Number of Successful Scans: 6245
Number of Failed Scans: 0
Number of Timedout Scans: 0
Number of Clean Files: 6240
Number of Infected Files: 5
Job Start Time: 9/2/2015 16:14:39
Job End Time: 9/2/2015 16:15:55

```
[[ID2d1c3692f1a0d2074b7116e0b2b1dba0]]
= vserver vscan scanner-pool apply-policy
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Apply scanner-policy to a scanner pool

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.
```

```
== Description

The `vserver vscan scanner-pool apply-policy` command applies a scanner
policy to the specified scanner pool on a specified Vserver.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
apply the scanner policy. The Vserver administrator can apply the scanner
policy to a scanner pool created within the scope of the Vserver or the
cluster. The cluster administrator can apply the scanner policy to a
scanner pool for any Vserver but cannot apply it within the scope of
cluster. The scope is determined at a Vserver level.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool.

``-scanner-policy <Scanner policy>`` - Scanner Policy::
This parameter specifies the scanner policy that you want to apply to the
specified scanner pool on a Vserver. Currently only system policies are
available. Available system policies are:
+
  * ``_primary_`` - Makes it active always.
  * ``_secondary_`` - Makes it active only when none of the primary
external virus-scanning servers are connected.
  * ``_idle_`` - Makes it inactive always.



``[-cluster <Cluster name>]`` - Cluster on Which Policy Is Applied::
This parameter specifies the name of the cluster on which you want to
apply the scanner policy of a scanner pool. By default, it is applied on
the local cluster. This parameter does not have any significance if the
cluster is not in a DR relationship.

== Examples

The following command applies a scanner policy to the specified scanner
pool on a specified Vserver.
```

cluster1::> vserver vscan scanner-pool apply-policy -vserver vs1
-scanner-pool p1 -scanner-policy primary -cluster cluster2

cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool p1
Vserver: vs1
Scanner Pool: p1

Applied Policy: primary
Current Status: on
Cluster on Which Policy Is Applied: cluster2
Scanner Pool Config Owner: vserver
List of IPs of Allowed Vscan Servers: 1.1.1.1, 2.2.2.2
List of Privileged Users: cifs\u1, cifs\u2

```
[[ID0a0e3323556bdd4072328188a8f7e9cd]]
= vserver vscan scanner-pool create
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Create a scanner pool

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool create` command creates a Vscan scanner
pool. Scanner pool is a set of attributes which are used to validate and
manage connection between ONTAP and external virus-scanning server, or
"Vscan server". It also specifies other parameters which are used for
connection management. After creating a scanner pool, a scanner-policy
must be applied to it using the command xref:{relative_path}vserver-vscan-
scanner-pool-apply-policy.html[vserver vscan scanner-pool apply-policy] .
The default applied policy is ``_idle_`` , which means the scanner pool is
inactive.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
create a scanner pool.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool. Scanner pool name
can be up to 256 characters long and is a string that can only contain any
combination of ASCII-range alphanumeric characters (a-z, A-Z, 0-9), "_",
"-" and ".".

``-hostnames <text>,...`` - List of Host Names of Allowed Vscan Servers::
This parameter specifies a list of host names or IP addresses of the Vscan
servers which are allowed to connect to ONTAP.
```

``-privileged-users <Privileged user>,...`` - List of Privileged Users::
This parameter specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name" and can be up to 256 characters long. Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.

``[-request-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Request Service Timeout [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a scan request. It specifies the time interval in which the node waits for a response from the Vscan server. If the timeout is reached, the node allows the file-operation if the applicable On-Access policy has scan-mandatory set to 'off'. If the policy has scan-mandatory set to 'on', then the node will retry the scan or disallow the file-operation depending on the remaining lifetime of the CIFS request. Valid values for this field are from 10s to 40s. However, if scan-mandatory is set to 'off', the effective value is limited to a maximum of 35s. The default value is 30s.

``[-scan-queue-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Scan Queue Timeout [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a scan request in scan-engine's queue. The value for this field must be between 10s and 30s. By default, it is 20s.

``[-session-setup-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Setup Timeout [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a response for session-setup-message. The value for this field must be between 5s and 10s. By default, it is 10s.

``[-session-teardown-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Teardown Timeout [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a response for session-teardown-message, or for any message to be received for a session-id, after the underlying connection has been disconnected. The value for this field must be between 5s and 10s. By default, it is 10s.

``[-max-session-setup-retries <integer>]`` - Max Number of Consecutive Session Setup Attempts [.advaced-priv]#(privilege: advanced)#::

```
This parameter specifies the maximum number of consecutive session-setup
attempts. The value for this field must be between 1 and 10. By default,
it is 5.

== Examples

The following example creates a scanner pool.
```

Cluster1::> vserver vscan scanner-pool create -vserver vs1 -scanner-pool SP
-hostnames 1.1.1.1,vmwin204-27.fsct.nb -privileged-users cifs\u1,cifs\u2

Cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool SP
Vserver: vs1
Scanner Pool: SP
Applied Policy: idle
Current Status: off
Cluster on Which Policy Is Applied: -
Scanner Pool Config Owner: vserver
List of IPs of Allowed Vscan Servers: 1.1.1.1, 10.72.204.27
List of Host Names of Allowed Vscan Servers: 1.1.1.1, vmwin204-27.fsct.nb
List of Privileged Users: cifs\u1, cifs\u2

```
== Related Links

* xref:{relative_path}vserver-vscan-scanner-pool-apply-policy.html[vserver
vscan scanner-pool apply-policy]




[[ID7a79646862c57430a66d6b3ccb0b51db]]
= vserver vscan scanner-pool delete
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Delete a scanner pool

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool delete` command deletes a scanner pool.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver from which you want to
delete a scanner pool.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner-pool that you want to
delete.

== Examples

The following example deletes a scanner pool.
```

cluster1::> vserver vscan scanner-pool delete -vserver vs1 -scanner-pool test

cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool test
There are no entries matching your query.

```
[[ID130bba56a6af2d000b3b53c47f8329d2]]
```

```
= vserver vscan scanner-pool modify
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Modify a scanner pool

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool modify` command modifies a Vscan scanner
pool. Scanner pool is a set of attributes which are used to validate and
manage connection between ONTAP and external virus-scanning server, or
"Vscan server". It also specifies other parameters which are used for
connection management.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver on which you want to
modify a scanner pool.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool. Scanner pool name
can be up to 256 characters long and is a string that can only contain any
combination of ASCII-range alphanumeric characters (a-z, A-Z, 0-9), "_",
"-" and ".".

``[-hostnames <text>,...]`` - List of Host Names of Allowed Vscan
Servers::
This parameter specifies a list of host names or IP addresses of the Vscan
servers which are allowed to connect to ONTAP.

``[-privileged-users <Privileged user>,...]`` - List of Privileged Users::
This parameter specifies a list of privileged users. A valid form of
privileged user-name is "domain-name\user-name" and can be up to 256
characters long. Privileged user-names are stored and treated as case-
insensitive strings. Virus scanners must use one of the registered
privileged users for connecting to ONTAP for exchanging virus-scanning
protocol messages and to access file for scanning, remedying and
quarantining operations.

``[-request-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
```

```
- Request Service Timeout [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a scan request. It
specifies the time interval in which the node waits for a response from
the Vscan server. If the timeout is reached, the node allows the file-
operation if the applicable On-Access policy has scan-mandatory set to
'off'. If the policy has scan-mandatory set to 'on', then the node will
retry the scan or disallow the file-operation depending on the remaining
lifetime of the CIFS request. Valid values for this field are from 10s to
40s. However, if scan-mandatory is set to 'off', the effective value is
limited to a maximum of 35s.

``[-scan-queue-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Scan Queue Timeout
[.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a scan request in scan-
engine's queue. The value for this field must be between 10s and 30s.

``[-session-setup-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Setup
Timeout [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a response for session-
setup-message. The value for this field must be between 5s and 10s.

``[-session-teardown-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Teardown
Timeout [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the timeout value for a response for session-
teardown-message, or for any message to be received for a session-id,
after the underlying connection has been disconnected. The value for this
field must be between 5s and 10s.

``[-max-session-setup-retries <integer>]`` - Max Number of Consecutive
Session Setup Attempts [.advaced-priv]#(privilege: advanced)#::
This parameter specifies the maximum number of consecutive session-setup
attempts. The value for this field must be between 1 and 10.

== Examples

The following example modifies a scanner pool.
```

Cluster1::> vserver vscan scanner-pool modify -vserver vs1 -scanner-pool SP
-hostnames 2.2.2.2,vmwin204-29.fsct.nb -privileged-users cifs\u3

Cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool SP
Vserver: vs1
Scanner Pool: SP
Applied Policy: idle

Current Status: off
Cluster on Which Policy Is Applied: -
Scanner Pool Config Owner: vserver
List of IPs of Allowed Vscan Servers: 2.2.2.2, 10.72.204.29
List of Host Names of Allowed Vscan Servers: 2.2.2.2, vmwin204-29.fsct.nb
List of Privileged Users: cifs\u3

```
[[ID233146d0fa74941fb8ef206c04ae34c3]]
= vserver vscan scanner-pool resolve-hostnames
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Resolve the hostnames configured in the scanner pool

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool resolve-hostnames` command resolves the
host names configured in the scanner pool and update it with the IP
addresses. This command also updates the active scanner pool configuration
of the Vserver if the scanner pool is part of that. You must run this
command for the scanner pool whose host name entry is modified in the DNS
server.

== Parameters

``-vserver <vserver>`` - Vserver::
This parameter specifies the name of the Vserver for which you want to
resolve host names.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool for which you want
to resolve host names.

== Examples

The following example resolves the host names of a scanner pool:
```

cluster1::> vserver vscan scanner-pool resolve-hostnames -vserver vs1 -scanner-pool SP

Cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool SP
Vserver: vs1

Scanner Pool: SP
Applied Policy: primary
Current Status: on
Cluster on Which Policy Is Applied: Cluster1
Scanner Pool Config Owner: vserver
List of IPs of Allowed Vscan Servers: 10.72.204.27, 10.72.204.29
List of Host Names of Allowed Vscan Servers: vmwin204-27.fsct.nb, vmwin204-29.fsct.nb
List of Privileged Users: cifs\u1, cifs\u2

```
[[IDd3cb1e0fe5913d27598a9672e1fdcf6b]]
= vserver vscan scanner-pool show-active
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display active scanner pools

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool show-active` command displays active
scanner pool information available to the Vserver. The active scanner pool
configuration is derived by merging the information of the scanner pools
which are currently active on a Vserver. If you do not specify any
parameters, the command displays the following information about all
Vservers:

* Vserver name
* List of scanner pools
* List of servers
* List of privileged user



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the specified Vserver.

``[-scanner-pools <Scanner pool>,...]`` - List of Enabled Scanner Pools::
If you specify this parameter, the command displays information only about
the Vservers that have the specified scanner pool or pools. A scanner pool
becomes part of this list if it is active at this time.

``[-servers <IP Address>,...]`` - Merged List of IPs of Allowed Vscan
Servers::
If you specify this parameter, the command displays information only about
the Vservers that have the specified server or servers. Servers of all
active scanner pools on a Vserver are merged to derive this effective
server list.

``[-privileged-users <Privileged user>,...]`` - Merged List of Privileged
Users::
If you specify this parameter, the command displays information only about
the Vservers that have the specified privileged user or users. Privileged
users of all active scanner pools on a Vserver are merged to derive this
effective privileged user list.

``[-request-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Request Service Timeout [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the Vservers that have the specified request-timeout. This is set to the
maximum value of the request-timeout of all active scanner pools on a
Vserver.

``[-scan-queue-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Scan Queue Timeout
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the Vservers that have the specified scan-queue-timeout. This is set to
the maximum value of the scan-queue-timeout of all active scanner pools on
a Vserver.

``[-session-setup-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Setup
Timeout [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the Vservers that have the specified session-setup-timeout. This is set to
the maximum value of the session-setup-timeout of all active scanner pools
```

```
on a Vserver.

``[-session-teardown-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Teardown
Timeout [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the Vservers that have the specified session-teardown-timeout. This is set
to the maximum value of the session-teardown-timeout of all active scanner
pools on a Vserver.

``[-max-session-setup-retries <integer>]`` - Max Number of Consecutive
Session Setup Attempts [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the Vservers that have the specified max-session-setup-retries. This is
set to the maximum number of the max-session-setup-retry of all active
scanner pools on a Vserver.

== Examples

The following example displays information about active scanner pool on
all Vservers.
```

cluster1::> vserver vscan scanner-pool show
Scanner Pool Privileged Scanner
Vserver Pool Owner Servers Users Policy
---------- ---------- ------- ------------------- --------------- ----------
Cluster clus cluster 5.5.5.5 cifs\u5 idle
vs1 new vserver 1.1.1.1, 2.2.2.2 cifs\u1 primary
vs1 clus cluster 5.5.5.5 cifs\u5 idle
vs1 p1 vserver 3.3.3.3 cifs\u4 primary
vs2 clus cluster 5.5.5.5 cifs\u5 primary
vs2 p2 vserver 3.3.3.3, 4.4.4.4 cifs\u2 primary
6 entries were displayed.

cluster1::> vserver vscan scanner-pool show-active
Privileged
Vserver Scanner Pools Servers Users
---------- ------------ ------------------------- -----------------------
vs1 new, p1 1.1.1.1, 2.2.2.2, 3.3.3.3 cifs\u1, cifs\u4
vs2 clus, p2 3.3.3.3, 4.4.4.4, 5.5.5.5 cifs\u2, cifs\u5
2 entries were displayed.

```
[[IDa99c613b3fbf33a41b2ddade65e31135]]
= vserver vscan scanner-pool show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
Display scanner pools

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool show` command displays information about
the Vscan scanner pools belonging to the Vserver. It also displays the
scanner policy applied to the scanner pool and its current status in
Vserver scope. If you do not specify any parameters, the command displays
the following information about all scanner pools:

* Vserver name
* Scanner pool
* Scanner pool owner
* Scanner policy
* Current status
* Cluster on which policy is applied
* List of servers
* List of host names
* List of privileged user



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the scanner pools for the specified Vserver.

``[-scanner-pool <Scanner pool>]`` - Scanner Pool::
If you specify this parameter, the command displays information only about
the specified scanner pool.

``[-scanner-policy <Scanner policy>]`` - Applied Policy::
```

```
If you specify this parameter, the command displays information only about
the scanner pools for the specified scanner policy.

``[-current-status {on|off}]`` - Current Status::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified status.

``[-cluster <Cluster name>]`` - Cluster on Which Policy Is Applied::
If you specify this parameter, the command displays information only about
the scanner pools that are applied to the specified cluster.

``[-owner <Configuration owner>]`` - Scanner Pool Config Owner::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified owner.

``[-servers <IP Address>,...]`` - List of IPs of Allowed Vscan Servers::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified IP address or IP addresses.

``[-hostnames <text>,...]`` - List of Host Names of Allowed Vscan
Servers::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified host name or host names.

``[-privileged-users <Privileged user>,...]`` - List of Privileged Users::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified privileged user or users.

``[-request-timeout <[<integer>d][<integer>h][<integer>m][<integer>s]>]``
- Request Service Timeout [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified request-timeout.

``[-scan-queue-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Scan Queue Timeout
[.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified scan-queue-timeout.

``[-session-setup-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Setup
Timeout [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified session-setup-timeout.

``[-session-teardown-timeout
<[<integer>d][<integer>h][<integer>m][<integer>s]>]`` - Session Teardown
```

```
Timeout [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified session-teardown-timeout.

``[-max-session-setup-retries <integer>]`` - Max Number of Consecutive
Session Setup Attempts [.advaced-priv]#(privilege: advanced)#::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified max-session-setup-retries.

== Examples

The following example displays information about all scanner pools.
```

Cluster1::> vserver vscan scanner-pool show
Scanner Pool Privileged Scanner
Vserver Pool Owner Servers Users Policy
---------- ---------- ------- ---------- ---------- ----------
vs1 SP vserver 1.1.1.1, cifs\u1, primary
10.72.204.27 cifs\u2
vs1 p1 vserver 3.3.3.3 cifs\u1, secondary
cifs\u2
2 entries were displayed.

```
The following example displays detailed information about one scanner
pool.
```

Cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool SP
Vserver: vs1
Scanner Pool: SP
Applied Policy: primary
Current Status: on
Cluster on Which Policy Is Applied: Cluster1
Scanner Pool Config Owner: vserver
List of IPs of Allowed Vscan Servers: 1.1.1.1, 10.72.204.27
List of Host Names of Allowed Vscan Servers: 1.1.1.1, vmwin204-27.fsct.nb
List of Privileged Users: cifs\u1, cifs\u2

```
[[ID085591fe14053f63e4259b2541242f02]]
= vserver vscan scanner-pool privileged-users add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to the list of privileged users

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool privileged-users add` command adds one
privileged users or list of privileged users to the specified scanner
pool.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
scanner pool on which you want to add a privileged user or users.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool to which you want to
add a privileged user or users.

``-privileged-users <Privileged user>,...`` - List of Privileged Users::
This parameter specifies a list of privileged users. A valid form of
privileged user-name is "domain-name\user-name" and can be up to 256
characters long. Privileged user-names are stored and treated as case-
insensitive strings. Virus scanners must use one of the registered
privileged users for connecting to ONTAP for exchanging virus-scanning
protocol messages and to access file for scanning, remedying and
quarantining operations.

== Examples

The following example adds a list of privileged users to the specified
scanner pool.
```

cluster1::> vserver vscan scanner-pool privileged-users add -vserver vs1
-scanner-pool p1 -privileged-users cifs\u2,cifs\u3

cluster1::> vserver vscan scanner-pool privileged-users show -vserver vs1
-scanner-pool p1
Vserver: vs1
Scanner Pool: p1
List of Privileged Users: cifs\u1, cifs\u2, cifs\u3

```
[[ID1a320635be9beff01bfc135dda6e3ad0]]
= vserver vscan scanner-pool privileged-users remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from the list of privileged users

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool privileged-users remove` command removes
one privileged users or list of privileged users from the specified
scanner pool. All the existing privileged users of a scanner pool cannot
be removed.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
scanner pool on which you want to remove a privileged user or users.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool from which you want
to remove a privileged user or users.

``-privileged-users <Privileged user>,...`` - List of Privileged Users::
This parameter specifies the privileged user or users that you want to
remove from the specified scanner pool.

== Examples

The following example removes a list of privileged users from the
specified scanner pool.
```

cluster1::> vserver vscan scanner-pool privileged-users remove -vserver vs1

-scanner-pool p1 -privileged-users cifs\u2,cifs\u3

cluster1::> vserver vscan scanner-pool privileged-users show -vserver vs1
-scanner-pool p1
Vserver: vs1
Scanner Pool: p1
List of Privileged Users: cifs\u1

```
[[ID5c9f60adfa06c0e514b22a6016e12c4c]]
= vserver vscan scanner-pool privileged-users show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of privileged users

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool privileged-users show` command displays
the list of privileged users of the Vscan scanner pools belonging to the
Vserver. If you do not specify any parameters, the command displays the
following information about the scanner pools:

* Vserver name
* Scanner pool
* List of privileged users



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
```

```
the scanner pools for the specified Vserver.

``[-scanner-pool <Scanner pool>]`` - Scanner Pool::
If you specify this parameter, the command displays information only for
the specified scanner pool.

``[-privileged-users <Privileged user>,...]`` - List of Privileged Users::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified privileged user or users.

== Examples

The following example displays the list of privileged users of all scanner
pools.
```

cluster1::> vserver vscan scanner-pool privileged-users show
Vserver Scanner Pool Privileged Users
--- --- ---
Cluster clus cifs\u5
vs1 new cifs\u7
vs1 clus cifs\u5
vs1 p1 cifs\u1, cifs\u2
vs2 clus cifs\u5
vs2 p2 cifs\u2
6 entries were displayed.

```
[[ID2521cf9010416fe6a49b9501032122d1]]
= vserver vscan scanner-pool servers add
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Add to the list of hostnames

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool servers add` command adds one server or
list of servers to the specified scanner pool.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
scanner pool on which you want to add a server or servers.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool to which you want to
add a server or servers.

``-hostnames <text>,...`` - List of Host Names for Vscan Servers::
This parameter specifies the host name or host names that you want to add
to the specified scanner pool.

== Examples

The following example adds a list of servers to the specified scanner
pool.
```

Cluster1::> vserver vscan scanner-pool servers add -vserver vs1
-scanner-pool SP -hostnames 2.2.2.2, vmwin204-27.fsct.nb

Cluster1::> vserver vscan scanner-pool servers show -vserver vs1 -scanner-pool SP
Vserver: vs1
Scanner Pool: SP
List of IPs of Allowed Vscan Servers: 1.1.1.1, 2.2.2.2, 10.72.204.27
List of Host Names of Allowed Vscan Servers: 1.1.1.1, 2.2.2.2,
vmwin204-27.fsct.nb

```
[[ID04ac0537952b8b6539ca5dc188375c1e]]
= vserver vscan scanner-pool servers remove
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Remove from the list of hostnames

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool servers remove` command removes one server
or list of servers from the specified scanner pool. All the existing
servers of a scanner pool cannot be removed.

== Parameters

``-vserver <vserver name>`` - Vserver::
This parameter specifies the name of the Vserver containing the specified
scanner pool on which you want to remove a server or servers.

``-scanner-pool <Scanner pool>`` - Scanner Pool::
This parameter specifies the name of the scanner pool from which you want
to remove a server or servers.

``-hostnames <text>,...`` - List of hostnames for Vscan Servers::
This parameter specifies the host name or host names that you want to
remove from the specified scanner pool.

== Examples

The following example removes a list of servers from the specified scanner
pool.
```

Cluster1::> vserver vscan scanner-pool servers remove -vserver vs1 -scanner-pool SP -hostnames vmwin204-27.fsct.nb

Cluster1::> vserver vscan scanner-pool servers show -vserver vs1 -scanner-pool SP
Vserver: vs1
Scanner Pool: SP
List of IPs of Allowed Vscan Servers: 1.1.1.1, 2.2.2.2
List of Host Names of Allowed Vscan Servers: 1.1.1.1, 2.2.2.2

```
[[IDe2f0416443a7b9bb7ed29351a27d50be]]
= vserver vscan scanner-pool servers show
:hardbreaks:
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Display list of servers

*Availability:* This command is available to _cluster_ and _Vserver_
administrators at the _admin_ privilege level.

== Description

The `vserver vscan scanner-pool servers show` command displays the list of
servers of the Vscan scanner pools belonging to the Vserver. If you do not
specify any parameters, the command displays the following information
about all scanner pools:

* Vserver name
* Scanner pool
* List of servers



== 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.

``[-vserver <vserver name>]`` - Vserver::
If you specify this parameter, the command displays information only about
the scanner pools for the specified Vserver.

``[-scanner-pool <Scanner pool>]`` - Scanner Pool::
If you specify this parameter, the command displays information only for
the specified scanner pool.

``[-servers <IP Address>,...]`` - List of IPs of Allowed Vscan Servers::
```

```
If you specify this parameter, the command displays information only about
the scanner pools that have the specified IP address or IP addresses.

``[-hostnames <text>,...]`` - List of Host Names of Allowed Vscan
Servers::
If you specify this parameter, the command displays information only about
the scanner pools that have the specified host name or host names.

== Examples

The following example displays the list of servers of all scanner pools.
```

cluster1::> vserver vscan scanner-pool servers show
Vserver Scanner Pool Servers
--- --- ---
vs1 SP 1.1.1.1, 10.72.204.27
vs2 p1 10.72.204.29
6 entries were displayed.

```
The following example displays the list of servers and host names of all
scanner pools.
```

cluster1::> vserver vscan scanner-pool servers show -instance
Vserver: vs1
Scanner Pool: SP
List of IPs of Allowed Vscan Servers: 1.1.1.1, 10.72.204.27
List of Host Names of Allowed Vscan Servers: 1.1.1.1, vmwin204-27.fsct.nb
Vserver: vs2
Scanner Pool: p1
List of IPs of Allowed Vscan Servers: 10.72.204.29
List of Host Names of Allowed Vscan Servers: vmwin204-29.fsct.nb
2 entries were displayed.

```
:leveloffset: -1


:leveloffset: -1


[[ID91c6790188552d1f9e47dd3a045d39d7]]
= Legal notices
:icons: font
:relative_path: ./
:imagesdir: {root_path}{relative_path}../media/

// Include the common notices
[.lead]
```

```
Legal notices provide access to copyright statements, trademarks, patents,
and more.

== Copyright

link:https://www.netapp.com/company/legal/copyright/[https://www.netapp.co
m/company/legal/copyright/^]

== Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks
page are trademarks of NetApp, Inc. Other company and product names may be
trademarks of their respective owners.

link:https://www.netapp.com/company/legal/trademarks/[https://www.netapp.c
om/company/legal/trademarks/^]

== Patents

A current list of NetApp owned patents can be found at:

link:https://www.netapp.com/pdf.html?item=/media/11887-
patentspage.pdf[https://www.netapp.com/pdf.html?item=/media/11887-
patentspage.pdf^]

== Privacy policy

link:https://www.netapp.com/company/legal/privacy-
policy/[https://www.netapp.com/company/legal/privacy-policy/^]

// Include the open source notice intro
== Open source

Notice files provide information about third-party copyright and licenses
used in NetApp software.

// Add required link(s)



:leveloffset: -1

<<<
*Copyright information*

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No
part of this document covered by copyright may be reproduced in any form
```

or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

```
*Trademark information*

NETAPP, the NETAPP logo, and the marks listed at
link:http://www.netapp.com/TM[http://www.netapp.com/TM^] are trademarks of
NetApp, Inc. Other company and product names may be trademarks of their
respective owners.
```